What I Wish I Had Known Early in Graduate School but Didn't — and How to Prepare For a Good Job Afterward

Sven Koenig

■ *Begin with the end in mind!*[1] *PhD students in artificial intelligence can start to prepare for their career after their PhD degree immediately when joining graduate school, and probably in many more ways than they think. To help them with that, I asked current PhD students and recent PhD computer-science graduates from the University of Southern California and my own PhD students to recount the important lessons they learned (perhaps too late) and added the advice of Nobel Prize and Turing Award winners and many other researchers (including my own reflections), to create this article.*

This article is a written version of a talk that I gave during the doctoral consortium of the International Joint Conference on Artificial Intelligence 2019 in Macau, China. When I was asked to give the talk, I wanted to give a talk about topics that PhD students care about, such as how to survive graduate school, how to pick good research topics, how to write good papers and give good talks, and how to find good jobs after graduate school. I decided to focus on what PhD students can do early in graduate school to prepare for life afterward, not what they should do around the times of their job applications (because other good guides have been written on how to apply and interview for jobs and negotiate good hiring packages). Because my last time on the job market was more than fifteen years ago, I sent a request for lessons learned to the University of Southern California's Computer Science PhD mailing list (that reaches both current PhD students as well as recent graduates) as well as my own students. I received lots of information that I integrated with my own advice and the advice of many other researchers, resulting first in my talk at the doctoral consortium and now in this article.[2]

 ISSN 0738-4602

General guidelines that apply to all topics mentioned throughout this article are to prepare early, work hard, and ask for advice but also learn by observing others and by doing and likely failing the first couple of times. You might also want to try to stand out in the crowd. There are many ways to do so (depending on your interests), for example, via awards, scholarships, winning contests, or founding start-ups. You might also want to apply for fellowships, which — besides being prestigious — could give you more freedom than working as a teaching or research assistant. But standing out in the crowd also means that you should not necessarily do what everyone else does. For example, currently lots of applicants to graduate school state "machine learning" as their intended research field, yet there are lots of faculty members that are interested in other artificial intelligence (AI) topics, at least in the United States. Thus, students who are interested in other AI topics have less competition and might have a higher chance of getting accepted to the universities of their choice. (Of course, you want to be honest about your interests to achieve a good match for you.) Finally, being able to say "no" can also be important. For example, you do not need to accept every job. But if you do accept one, then you should try to do a good job and leave a good impression on people.

To find a good job after graduate school, you first need to understand what potential employers are looking for as well as your own interests and strengths and then obtain the right skills. I will discuss skills related to doing research, teaching, marketing, and managing your advisor. Marketing skills include writing good papers, giving good talks, and networking. The title of this article already suggests that I believe that marketing skills are important because you want to make sure that others know about you and your research, otherwise your research will not have its full impact and you might be invited for only a few job interviews. You thus want to invest a lot of effort before and during your job search into finding people who are interested in you and your research. You will hear "no" a lot, but it only takes a "yes" or two to get a good job. I will cover these skills briefly because this topic has already been widely discussed.

Of course, you need to have not only the right skills but also evidence that you possess them. After all, potential employers are looking for evidence that you did not exaggerate your skills in your job application, which is why portfolios (i.e., compilations of materials that provide evidence of your skills) and references (i.e., people who can attest to your skills) are important. For example, a portfolio to showcase your teaching skills might include teaching assistant awards that you have won, homework assignments that you have created, and slides and videos of lectures that you have given. Your job application typically does not include such portfolios directly but rather a pointer to a webpage with them. Interviewers and hiring managers will look for them.

Understanding Potential Employers[3]

Most potential employers will ask themselves the following questions: Do we want to spend the next five to twenty-five years with you? Are you complementary with our existing capabilities? Do you have the right expertise and flexibility for the future? Are you a self-starter and able to work without supervision? Are you motivated and energetic? Would you accept our offer?

Most of these questions are self-explanatory. For example, the first question implies that your personality is important and you want to demonstrate your best behavior. However, the last question might need an explanation. When an employer decides to make you an offer, they will start to invest resources in you. For example, they will negotiate a hiring package with you with lots of perks (such as a relocation allowance and, in industry, a sign-on bonus) and might pay for you to travel to their location again, this time with your family, so that you can check it out thoroughly. If you decline their offer, they will have lost time and money. Thus, employers might not make an offer to their best applicant if they believe that the chance of getting them to accept is low. Thus, you want to avoid giving the impression that you sent the exact same application to many potential employers. For example, you might want to mention in the letter that accompanies your job application why you want to work for this specific employer. If you have friends or family at their location, you might want to mention that as well.

Potential employers will also ask additional questions, depending on their type of organization. *Research universities* (or research institutes) will ask themselves the following additional questions: Can you manage a research group? Will you collaborate with other researchers, especially those employed by us? Will you bring in funding and students? Do you have a long-term research vision that you intend to pursue? Will you become visible for your research?

The second question might need an explanation. A lot of research these days is done in groups, so it is important that you are interested in collaborating with other researchers, and able to do it. Of course, you want not only to work with researchers at other organizations but also to use the resources provided by your own organization, which will also help you to obtain more and larger grants for it. Another reason for trying to build bridges to all researchers that you meet during your job interview, not just the ones in your own research field (where building bridges should be easy for you), is that politics influences decision-making in any organization, including research universities: For example, researchers often cannot agree who to hire, especially in cases where only one position is available and its research field

has been left unspecified. Then, it often happens that the theory researchers try to hire a theory applicant, while the systems researchers try to hire a systems applicant. The AI researchers might try to hire you as an AI applicant. If the theory researchers do not like the systems applicant and vice versa and you have managed to build bridges with everyone, then you might be second on the hiring lists of both the theory and systems researchers (because you would enhance their research), and all of them can easily agree that you would be a good applicant to hire. Thus, you want to research everyone who will talk to you during your job interview. You can then discuss joint research possibilities with them, which allows them to discuss their research with you, you to discuss your research with them, and each of you to make plans for the future, which also demonstrates your interest in working at that organization.

The fourth question might also need an explanation. During a job interview, you might get asked where you and your research field are going in the next five to ten years, so you might want to think about this early and form your own opinion. For example, we currently live in the era of big data, where our online and other activities result in lots of collected data, and machine learning is used to obtain models that help us to make sense of the data. However, ultimately one wants to use the models to make good decisions. Indeed, one often uses deep reinforcement learning these days to learn mappings from the input directly to actions. However, we really need to interface machine learning in the future much better with planning and other decision-making technologies to be able to make good decisions with reasonably small amounts of data and with the ability to explain the decisions, resulting in an era of big decisions rather than big data. Of course, you want to come up with your own vision of the future, for example, while attending research talks.

If you are interested in working at a research or teaching university, you might want to go to job talks of faculty members that interview at your university because you will have to give similar talks during your job interviews.

Teaching universities often ask themselves very different questions than research universities: Are you ready and committed to teaching our students?

Teaching universities place more emphasis on teaching and less on research. You are under less pressure to bring in funding for your research than at research universities, but you have to teach more and might not have many research-oriented students available to help you with your research.

If you are interested in working at a teaching university, you might want to be a teaching assistant several times (to acquire teaching expertise) and put together a strong teaching portfolio. Often, teaching universities even look for applicants that have taken responsibility for teaching at least one full course because the applicants then understand how time-consuming and demanding teaching is, and are more likely to be effective immediately after they get hired. Also, at least one letter of recommendation should address your teaching skills. One of my PhD students wanted to work at a teaching university, and initially did not get hired by the teaching universities of his choice because he graduated from a top research university. A lot of teaching universities suspected that he really wanted to do research but had not been accepted by a research university, and thus had decided to park at a teaching university without really being interested in teaching and with the intent to reapply to research universities next year. So, it would not make sense to offer him a job. This shows how carefully you need to prepare. He eventually took a temporary sabbatical substitute position at a teaching university for a year and then reapplied to the teaching universities of his choice. Now being a credible applicant, he got his dream job.

Industry will ask themselves the following additional questions: Can you present to and convince clients? Will you do applied work without prodding? Do we have current or upcoming projects for you? Will your skills or connections help us win new contracts?

If you are interested in working in research and development in industry, you might want to do one or more internships in industry or perhaps maintain a larger piece of software, either for your research group or as a hobby.

Postdoctoral positions, different from the aforementioned positions, are meant for temporary employment, often at research universities (but now more and more also in industry). The hiring researchers will ask themselves the following additional questions: Will you add complementary new ideas to our research group? Will you work without supervision? Can you manage an ongoing or upcoming project?

Many postdoctoral positions are not advertised and thus might require you to ask around and contact research group leaders in your research field directly, at least in the United States. The job responsibilities of postdoctoral researchers vary, so it is important to inquire about them. Because it is understood that postdoctoral positions are meant for temporary employment, it makes sense to keep applying for long-term employment every year while you are occupying a postdoctoral position. In fact, you might want to apply for postdoctoral positions early during your job search. When offered one, the hiring researcher will often (but not always) be willing to wait while you continue your job search, which means that you have a backup job and thus can go with confidence into the remainder of your job search.

Understanding Yourself

Now that you understand the questions that you might want to address directly or indirectly in your job applications and during your job interviews, you will also need to find out more about yourself. This is important because your research in graduate school

will hopefully influence what you do for the rest of your life. Graduate school is long, and your life is even longer. Thus, you should be in the driver's seat and cater to your interests, which requires you to ask yourself the following questions: What is your motivation for attending graduate school? What kind of work do you want to do afterward? Which research questions are you interested in? Which skills do you have or would you like to acquire?

You might have to experiment a bit to be able to answer these questions. For example, I always recommend to my PhD students to do one or more internships during their first two summers before starting to dedicate most of their time to research.

The first question might need an explanation. For example, money is probably not a good motivation for attending graduate school because you will not make much money during graduate school, at least in the United States. Of course, you need a PhD degree for some jobs — but you do not for others. For example, your higher salary in industry after graduate school might not make up for the time you spent in graduate school because those who joined industry earlier have already received several promotions by the time you graduate. The first question is thus especially important if you attend graduate school after having worked in industry because your old friends then have money to spend in their spare time while you might not.

The third and fourth questions might need a longer explanation. As a PhD student, you have probably done some research already when you are reading this article, but — it turns out — thinking about the research questions that you should be working on is still important for you because finding a good research topic can be a long process.

AI researcher Alan Newell (who received a Turing award in 1975) is quoted as having said:

> The scientific problem chooses you; you don't choose it. My style is to deal with a single problem, namely the nature of the human mind. That is the one problem that I have cared about throughout my scientific career, and it will last me all the way to the end.[4]

In other words, finding a good research topic is like a marriage. You meet lots of people in your life, might experiment by dating some of them, (ideally) eventually find the perfect match for you, and get married. Similarly, you will eventually find a research topic that is a great match with your interests and skills. This also means that you should not choose to do research on some topic just because everyone else seems to be working on it or because there seem to be lots of jobs currently available in that research field. While Alan Newell could express his research topic very concisely, I still need more text to describe mine. I always worked on multiple research topics that seemed to be unconnected, and it took me several years to understand what they had in common. I typically explain these days that I develop techniques for planning for both single and multiple agents that work well even if the agents have incomplete knowledge of their environments, imperfect abilities to manipulate them, limited or noisy perception, or insufficient reasoning speeds. A lot of my research currently addresses the last limitation, namely that planning can be slow, but decisions often need to be made in real-time (e.g., to avoid robots being idle). Thus, I work on exploiting domain structure to speed up decision-making. Being able to characterize your research topic concisely also has the advantage that you can use it as the short version of your elevator talk, so read on!

AI researcher Herb Simon (who was the advisor of Alan Newell and received a Turing award together with him in 1975 but also a Nobel prize in economics in 1978) is quoted as having said:

> I advise my graduate students to pick a research problem that is important (so that it will matter if it is solved) …[5]

In other words, for your research to have impact some people must care about your results. The more people care, the larger your impact. You do not need to pick *"P = NP?"* as your research topic but rather want to find your own research topic based on your vision of what is or will be important in your research field. You want to be bold in your choice initially. You can then narrow it down progressively to make it more doable until you arrive at a dissertation topic. While attending research talks, you might thus not only want to think about where your research field is going in the next five to ten years but also about important research topics — and do not forget to continue identifying new important research topics from time to time even when you are already working on some research topic.

Herb Simon continued:

> … but one for which they have a secret weapon that gives some prospect of success. Why a secret weapon? Because if the problem is important, other researchers as intelligent as my students will be trying to solve it; my students are likely to come in first only by having access to some knowledge or research methods the others do not have.[5]

In other words, you are essentially in a friendly(!) competition with other researchers in your research field, so you might want to exploit any competitive advantage that you have over them. This could be a skill (for example, you might be able to develop complex software or prove complex properties of your algorithms), scientific knowledge (for example, you might know about control theory in addition to AI), or application knowledge (for example, you might know about oil drilling in addition to AI). You can then focus your research effort by picking research topics that benefit from your secret weapon.

It took me several years to understand what my secret weapon is, but I eventually realized that I often combine insights from different research fields when developing decision-making algorithms. I sometimes

do this myself, but I also often collaborate with experts in research fields other than AI. Collaborative research efforts require skills, and not all researchers find them easy to pursue. When you work on a bigger research problem (or application), you need to isolate a part of it that is better tackled by other researchers. Then, you need to find one or more suitable experts in the right research fields and describe that part of the research problem (or, more likely, a slight abstraction) to them in their language, so that they both want to solve it and are able to solve it. Finally, their solution must help you to solve your bigger research problem.

As an example, consider how auctions had been used in robotics to assign tasks to robots. Robotics researchers had realized that auctions might be an effective and practical approach to the coordination of robot teams. The robots of a robot team bid their costs on tasks, and the robot with the lowest bid on a task wins the task because it is most suitable for it. Such auctions are communication-efficient because information is compressed into bids. They are also computation-efficient because the robots calculate their bids in parallel. Finally, they are also robust if they can be implemented in a decentralized way because the failure of one or more robots then only degrades the quality of the resulting robot-task allocation rather than preventing the robot team from functioning. Thus, robotics researchers had experimented with different auction mechanisms and noticed that some of them work well and others do not — but it had remained unclear why this was the case or how to design better auction mechanisms. As can be expected, if one understands AI and robotics on one hand and economics on the other, one has a secret weapon over researchers that understand only one of these two research fields because one now stands on the shoulders of two giants. For example, one can take inspiration from the different auction mechanisms that economics researchers have developed and studied. However, the setting in economics is typically a competitive setting, where the participants are self-interested. Thus, they have concerns about other participants gaming the system (for example, via collusion or shilling) and thus about revealing their own preferences. The setting in robotics is typically a cooperative setting, where all robots are programmed to maximize the utility of their robot team, and there is pressure to clear the auctions in real-time to avoid the robots being idle. These differences imply that insights from economics cannot be applied in a straightforward way to robotics. A great deal of research is required to account for them. The tasks in robotics are often navigation tasks, where robots must move to given locations such that the sum of their travel times is small. For example, National Aeronautics and Space Administration scientists might be interested in the composition of several rocks on Mars. In this case, each rock must be visited by one robot so that a rock probe can be taken from it. The robots thus bid on the rocks and then find shortest paths that visit all rocks that they have won. One can exploit domain structure to achieve a high-quality robot-rock allocation in real-time, for example, by realizing that robots can amortize their travel times when visiting several nearby rocks each. The problem that the robots need to solve is a decentralized version of a vehicle-routing problem from operations research, where trucks need to be routed to deliver goods to customers. This insight helps one to exploit the domain structure and implies that, if one understands AI and robotics, economics, and operations research, one has an even larger competitive advantage. Thus, an interdisciplinary team of experts is a secret weapon because there are not many individual researchers or teams of researchers that know all three research fields.[6] In general, I believe that there should be a science of making intelligent decisions that combines insights from several research fields, including AI, operations research, decision theory, economics, control theory, statistics, theoretical computer science, and other research fields. This way, researchers would have larger toolboxes available to build intelligent systems, likely resulting in smarter agents and robots.[7] Thus, I regularly suggest that my PhD students take not only AI classes but also classes in other decision-making fields. This is not a new insight. For example, the most popular textbook in AI already incorporates content from different research fields, including utility theory and multiattribute utility functions from decision theory, game theory and auctions from economics, and stochastic dynamic programming from operations research.[8]

The idea of using your secret weapon(s) also extends to areas other than your research. For example, if you are good with organizing events, you could stand out in the crowd by organizing a workshop at a conference. If you are good in teaching, you could give a tutorial on your research at a conference. Some advanced PhD students co-organize workshops and tutorials at conferences, typically by partnering with other researchers.

Obtaining the Right Skills

I will now discuss the various skills that you might want to acquire in graduate school.

Research

One fundamental research skill is being able to solve research problems. A second, related fundamental research skill is being able to identify good research problems and break them down into doable pieces. The latter research skill is as important as the former for independent researchers. Thus, you want to practice this skill in graduate school by not only solving research problems given to you by others, although it is fine to ask your advisor for interesting research problems initially. Practicing this skill also ensures that you cater to your interests and skills, and work on the research problems that you really want to

work on. Advisors often want you to work on your own research problems as well (although not always, for example, if they are under pressure to demonstrate progress on their funded research projects). For example, I often use PhD students to branch out into new research problems, which is almost impossible for me to do otherwise, because I often have too little time left for my own research. Also, students who work on their own research problems are often very productive because they are motivated to work on them and might have secret weapons for solving them. My PhD students are also smarter than me and have better research ideas.

Graduate school exists to ensure that you acquire these fundamental research skills.[9] Thus, I will mention only four pieces of advice: You should not isolate yourself by always working from home or in your office with your door closed, although some quiet time where you can completely concentrate on your research is often helpful (which might not be at home with your TV running); you should keep a notebook with your research questions, ideas, and insights; you should avoid aimlessly reading the literature because then everything will appear to have been solved already (rather, you should pick a research problem to work on and then, in parallel, work on it and read the relevant literature, and suddenly you will notice that very few research problems have been solved already and that the published solutions often have disadvantages or flaws); and you should not forget to archive the results of your research, which include not only your papers but also your software, data, and notes. You want to archive the exact software that you have used to produce the data that you have reported in your papers, perhaps in addition to a cleaned-up version (in which you might have inadvertently introduced bugs) because it might take several years after publication before you get approached with questions about it. For example, one of my PhD students had used the 8-puzzle to demonstrate the power of his new heuristic search algorithm. He got approached after a year or two by a researcher who could not reproduce his results. A quick check of the archived software revealed that my PhD student had used the goal configuration in figure 1A. He had not described it in his paper because he considered it to be the standard goal configuration. The other researcher had assumed that the goal configuration was the one on the right in figure 1B. Once this issue was clarified, the other researcher was able to reproduce the reported results. It would have been difficult to resolve this misunderstanding without the archived software.

A third research skill is to always be critical. A PhD student stated:

> "Never believe others — not what you read and not your advisor. Common wisdom quite often turns out to be wrong! Researchers need to be critical of common wisdom. Papers often use case studies on examples that work well for their research, and authors are often blind to the issues of their research or do not publish the disadvantages of it. Checking the results of published research is therefore often helpful."

A fourth research skill is to always consider how to design and apply your AI technology to benefit society the most and identify potential ethical issues related to it. For example, you might want to ask yourself how to guarantee the reliability, robustness, and safety of systems built with your technology; when to provide oversight of their operations; how to guarantee that their behaviors are consistent with social norms and human values; and how to ensure that they impact the standard of living, the distribution and quality of work, and other social and economic aspects in the best possible ways.[10] Keeping these questions always in mind as you do your research might not only avoid undesirable effects of your research but also result in interesting research topics for you to pursue.

A fifth research skill is never to lose sight of the overall research problem and be broadly interested, even beyond your research field. For example, you might want to go to colloquia at your university (including those outside of your research field or even computer science), take classes outside of your research field, or read interesting books outside of your research field. You might also want to attend tutorials or talks at conferences on research topics that you know little to nothing about but which sound interesting to you. This way, you might learn something that gives you unexpected insights into your research problem or unexpected applications of your technology in new research fields, in other words, obtain a secret weapon. Also, you cannot work on the research topic of your dissertation forever after graduation. Although depth is typically considered to be much more important for a dissertation than breadth, you need flexibility to work on new research topics after your graduation (for example, because your job now requires it). At that point, your breadth of knowledge will come in handy. Even where you have a large influence on your research topics, for example, if you become a professor at a research university, having broadened your research interests beyond the research topic of your dissertation is an important consideration for giving you tenure, at least in the United States, simply because it is deemed necessary for you doing cutting-edge research in the long run.

At the same time, you need to stay focused on completing graduate school, so cannot afford too many distractions. Good time management is important. For example, you might want to set yourself an ambitious but reasonable time frame for the completion of your PhD degree, perhaps five years in the United States. Of course, there are many uncontrollable influences, and thus you do not want to get discouraged if you feel that you are not as lucky as some of your co-students because not everything goes smoothly for you. By luck, I mean, for example,

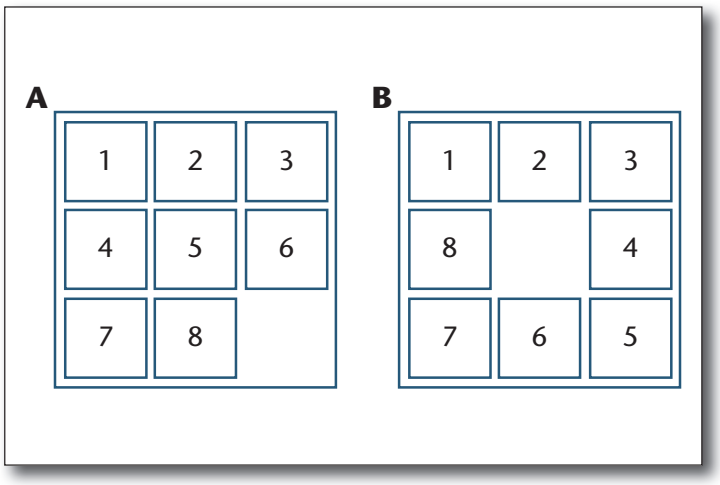

Figure 1. Goal Configurations of the 8-Puzzle.

that their first conference or journal submission was received well by the reviewers (and it boosted their confidence) or a senior researcher pointed them in a new research direction (and it turned out to be promising). You can boost your luck by discussing your research and the obstacles that you encounter with others. It also helps if you are passionate about your research and being in graduate school and do not constantly compare yourself to other PhD students. Do not be your own worst critic and stress yourself out. Ultimately, as always in life, it is up to you to overcome any obstacles in your way and make use of the help that is available to you. Remember that you were accepted to graduate school because people believed that you would succeed. For a few PhD students, success might be easy to achieve, but the success of most PhD students is proportional to the time that they spend on their research.

In general, you might always want to have a main objective to focus on, which you can change as needed. Diversions can be good (for example, to explore your research interests and acquire breadth), but Alan Newell is quoted as having said: "Diversions occur, make them count; salvage what is possible for the main goal."[11] In other words, you need to ensure that you get something out of them for finding a good job after graduate school. For example, your advisor might ask you to do a side job for the benefit of your research group or project. Remember that you do not need to accept every job, although you probably want to be a good citizen from time to time and help your advisor and your research group because they help you as well. If you accept the job, then you should try to do a good job at it. But you might want to follow Alan Newell's advice and make it count. For example, it might help you to ensure your emotional well-being if it results (often different from research) in immediate gratification, help you learn more about yourself, or help you assemble a portfolio. A colleague also pointed out that diversions can have long-term benefits that are not immediately obvious. For example, you should not hesitate to pursue a research direction for a while just because it intrigues you, even if you do not see any immediate benefits (for example, because you do not obtain immediate results or cannot relate it to your dissertation research). A lot of dissertation-relevant or impactful research has started as intriguing diversions. This is my experience as well. I worked on different research directions during graduate school. When I got stuck with my main research, I would often look at one of the other research directions during my breaks and, because there was no pressure on me to make progress on it, I could be relaxed and playful, which helped me to generate new ideas. I then switched to this research direction, only to get stuck after making some progress, and the cycle repeated. So, having more than one research direction helped me to make progress on my research.

A sixth research skill is to be able to work as part of interdisciplinary research groups. In graduate school, working as part of interdisciplinary research groups can help your research, as described earlier. After graduation, it can help you because you will often have to work as a part of a team with people whose backgrounds are very different from your own background, whether you still do research or not.

A seventh research skill is to be able to mentor other students. In graduate school, mentoring other students and helping them to solve their research problems can help you because it gives you experience in doing research with others and makes you think about research problems that you might not think about otherwise. It might also result in you becoming a coauthor of papers that other students are mainly responsible for. After graduation, it can help you because you will often have to mentor other people, whether you still do research or not.

An eighth research skill is to be able to acknowledge your mistakes. You always want to work carefully in your research, but you will make a mistake at some point in time. It is important for you to admit it and correct it because other researchers might otherwise build their research on your flawed research results, which will impede their scientific progress. You might worry about other researchers then thinking less of your ability to solve research problems. You will accordingly need to remember that every researcher makes mistakes, not just you, and that good researchers admit and correct their mistakes. I was once part of a research team that developed a new heuristic search algorithm and evaluated it in a testbed that we had developed. The algorithm beat the competition, and so we published the results — only to notice shortly afterward that we had made an implementation mistake in the testbed. Once corrected, our algorithm was no longer the best. We updated our paper, and I put — together with the correction — a warning on my website that the paper was no longer worth reading (you can still find it there). That was about ten years ago, and I am not aware of any disadvantage that

I had because of the correction. Today, different from back then, many digital archives have followed journals in publishing such errata. There is no stigma attached to publishing errata if every researcher does it when necessary, so I recommend that you do that as well.

Writing Papers and Dissertations[12]

You should submit your research to top AI conferences whenever possible because the quality of your research output counts more than its quantity in our research field. Top conferences in AI are much more tightly reviewed than other conferences, and the acceptance of your research to top conferences thus speaks to its quality. Many AI researchers consider publication in a top AI conference to be equivalent to publication in a top AI journal. Still, you might want to publish in journals as well because not everyone shares this opinion; for example, because they consider journal publications to be more comprehensive than conference publications due to their larger or missing page limits, or because their backgrounds are in research fields where journal publication is more valued than in AI.

To prepare for writing good papers, you should ask your advisor for advice on what makes a good paper because, after all, they need to be happy with the work that you write and publish. You should also read and criticize writing by others (ideally in a group because this often results in more insights) and review conference and journal submissions. Your advisor might be able to give you such submissions for review. Almost all authors think that their writing is easily understandable by others, so it is often a revelation when you read submissions and realize that the authors must have thought the same even though you see lots of issues with them. This experience helps you to avoid similar issues when you write your own papers. It also helps you to understand how reviewers think and thus how to write papers that they are willing to accept.

When writing for submission, you need to know your audience. For example, when you send a paper to a venue that you are unfamiliar with, you first want to look at publications in that venue to become familiar with how its articles and papers are written. You should let others comment on your manuscripts before you submit them. Your advisor will likely do that, but you might also want to make a deal with another PhD student in your research group to review each other's manuscripts. Even better is a deal with a PhD student from another university who works on a similar research topic because that student likely has no additional knowledge of your research and thus is in the exact same position as the reviewers. But you should give them a version of the manuscript to review that you consider to be the final version because comments on poorly written manuscripts typically focus on surface problems and thus will not be helpful to you.

You should not be discouraged if one of your submissions gets rejected because there are just too many uncontrollable influences in the review process, starting with the selection of the reviewers. If your submission gets rejected, you want to address all comments of the reviewers before submitting it again. Our research community is sufficiently small that there is a chance that you will get the same reviewer for the next submission even if it is to a different venue, and they would certainly notice that you did not take their comments into account. If your paper gets accepted, you also want to address all comments of the reviewers before submitting the camera-ready version. They might look at the final publication and would certainly notice that you did not take their comments into account. Either way, your future submissions might be at a disadvantage. You should never dismiss comments of reviewers easily even if they are wrong, which is advice that I received from my own advisor. For example, a reviewer might claim that you forgot to define an abbreviation even though you did define it. Rather than dismissing the comment as wrong, you should interpret it as meaning that they noticed the abbreviation, tried to find its definition, and failed. Consequently, you should highlight its definition more.

One of the most important pieces of advice on writing is to have a clear and concise hypothesis.[13] A hypothesis, also called a *claim*, is a statement that can be proven to be true or false. It is also the key to writing a good dissertation (which is why a dissertation is also often called a PhD thesis) because it focuses the dissertation much better than the phrase "This dissertation focuses on …." I learned this only after writing my own dissertation. A dissertation should not contain the contents of all your publications on your research topic or all thoughts that you have ever had on it but rather only those that help you to prove or disprove your hypothesis. For example, my PhD student William Yeoh, now a professor himself, wrote a dissertation with the title "Speeding up Distributed Constraint Optimization Search Algorithms." The first couple of sentences in the abstract (and the main text) of his dissertation read as follows:

> Distributed constraint optimization (DCOP) is a model where several agents coordinate with each other to take on values to minimize the sum of the resulting constraint costs... This model is becoming popular for formulating and solving agent-coordination problems. … DCOP search algorithms can be viewed as distributed versions of centralized search algorithms. Therefore, I hypothesize that one can speed up DCOP search algorithms by applying insights gained from centralized search algorithms, specifically (1) by using an appropriate search strategy, (2) by sacrificing solution optimality, (3) by using more memory, and (4) by reusing information from searches of similar DCOP problems.[14]

The first sentence explains the topic and some of the terminology used in the title. The second sentence explains the importance of the dissertation. The third

sentence states the main insight that enabled the research described in the dissertation and explains more of the terminology used in the title. The fourth sentence is the hypothesis. In William Yeoh's dissertation, it is followed by a fifth sentence that explains how he proves this hypothesis in the rest of his dissertation.

The structure of a paper is simple. In the introduction, you describe the research problem, your motivation for studying it (for example, its importance), and your contributions (typically your hypothesis with a short outlook on the main results). In the main part, you prove your hypothesis, for example, by developing an algorithm, analyzing its properties, and experimentally evaluating it. In the conclusions, you summarize your hypothesis and results. This means that you convey your insights three times to your audience. You first tell them what you are going to tell them, then you tell them, and finally you tell them what you told them. You should write the introduction and conclusions especially carefully because many readers will read only them and not the main part of your paper (for example, because they only want to get a general idea about its topic and insights).

Giving Talks

Once your conference submission has been accepted and you are preparing the conference presentation, you need to think about its structure, the slides, and the talk. You should use your presentation mainly to advertise your research and thus, as Rao Kambhampati puts it, stress the *why* and *what* more than the *how*. Other than that, the structure of your presentation can follow the structure of your paper. Your slides should make this structure clear to the audience and remind it before each part of your presentation where you are in your presentation. I typically first describe the research problem and its motivation. I then present an executive summary slide that contains the take-home message of the presentation in plain English and point out that the audience should remember this one slide about the talk. The take-home message of the presentation is typically related to your hypothesis and results in the paper. I next present the technical part of the presentation and conclude with a summary slide that is like the executive summary slide but often contains more details and perhaps also future work. I leave this slide projected when answering questions so that the audience remembers the gist of the presentation during this time and hopefully also long afterward. Ending on a slide that simply states "Questions?" is counterproductive.

You want to make your presentation easy to follow; for example, use examples rather than formalisms and use plain English rather than unusual terminology or abbreviations. It is hard for an audience to remember more than one new term or abbreviation and, if you really need to use one, you want to emphasize it during its introduction so that your audience knows to remember it. You might want to summarize each part of your presentation, especially technical parts where you might lose part of your audience. A short summary allows it to catch up and thus re-engages it. Your audience will be busy listening to you and therefore does not have the cognitive capacity to comprehend complex slides or draw any conclusions on their own. Thus, you do not want to use PowerPoint-style animations but instead tell your audience exactly what to think (which is also the reason for the executive summary slide) and keep the amount of text and formulas on the slides small. Ideally, you want to convey one idea per slide. You do not need to put everything that you want to say on your slides. Rather, your slides should support what you say with illustrations that you can point to during your talk. You should use a large font size so that your slides can be read easily. Also, you should acknowledge prior work to show that you are an expert and please the people you mention.

You should not read your talk from a script, but you should vary your voice to avoid people falling asleep during your talk, highlight important facts, and express enthusiasm. After all, if you are not enthusiastic about your own research, then why should the audience care about it? And you should get the length of your talk right, which is difficult to do without a practice talk.

Once you have practiced your talk by yourself, you might want to give a practice talk to your research group and get feedback from others. You might also want to tape your practice and actual talks and listen to the recordings afterward.

Teaching and Mentoring

I suggest that you be a teaching assistant for one or two classes, whether your university requires it or not, because this allows you to practice skills that are important for life after graduation, whether you want to work at a university or not. Both lecturing and mentoring help you to explain complicated matters well because students provide you with feedback about what they understand (much more than, say, conference audiences). For example, lecturing allows you to practice how to tailor your presentations to your audience. It also helps you to gain breadth and depth because, as they say, you do not really understand any material until you have taught it and had to answer questions about it.

Networking

A PhD student stated: "Who we know may be more important than what we know. The people and friends I made in graduate school have significantly shaped my career trajectory. So, step outside of your comfort zone, talk to peers, strangers, professors. You never know who you will meet." For example, your contacts might be helpful for research collaborations. In the context of finding a good job after graduate school, your contacts might be helpful for several reasons: Some of them might want to hire you; some of them might be able to provide you with sample job applications; some of them might be able to point out job opportunities to you; some of them

might spread the word for you that you are looking for a job; and some of them might be able to vouch for you. This last group of people is extremely important. You will need references for your job search, especially from experts on your research topic who know you well.

You want to network constantly. For example, you can meet people in your department by talking to all faculty members (not just your advisor) and visitors. You might be able to sign up for talking to short-term visitors after their talks, and you can visit longer-term visitors in their offices. Often, they do not know many people in the department yet and thus will be especially happy to talk to you. You can meet people at conferences by not hanging out with your own research group. Instead, you want to attend a doctoral consortia and social events, go to lunch with people you just met at the conferences, and approach people who share your research interests after talks and at posters (for example, by asking questions about their research). You can also invite people to your own talks and posters. Recently, a student interrupted a conversation that I had with a colleague in the far corner of the room where the student presented a poster. The student pointed out that she had built on my research and insisted that I listen to her poster presentation. I did, enjoyed learning about her research, and had a longer conversation with her afterward, which resulted in an e-mail exchange about related research. You do not need to be afraid of approaching senior researchers. Most people love to talk about their research, so you can mention that you have read their publications and use their ideas in your research, which also gives you an opportunity to talk about your own research. Of course, you can also try to do that by e-mail (for example, by sending them your publications), but approaching them in person is often more successful. Do not forget to introduce your friends to your contacts, and vice versa.

The more a contact knows about you, the more helpful they are as a reference. Your advisor knows you best and should be one of your references, but they are expected to say good things about you, which is why their opinion is often somewhat discounted. Thus, you might want to work with other experts in your research field closely. For example, you might want to start joint research with visitors to your research group or with research mentors at other universities. You might want to do internships in industry or visit other research groups (including at universities) for longer periods of time. You might also want to have a dissertation committee member from another university who is an expert on your research topic if your university allows it. However, never assume that someone is willing to be your reference before you asked them and they have agreed.

When you meet new researchers, they will often ask you about your research. You should have two replies ready, namely one that is about two to three sentences long and one that is about two to three minutes long. Such replies are also important for job interviews and are often called "elevator talks" because someone who had to miss your job talk might approach you, for example in an elevator, with the question of what your research is all about and the circumstances then require that you have a concise but informative answer handy. In addition, you also want to be able to explain the importance of your research and its expected (or actual) impact on your research field and society.

You should also have a well-designed website with information about you, including your Curriculum Vitae. You can also put your publications, software, and data on your website because their impact will be larger if they are easy to find. I add abstracts to the papers on my website in the hope that this makes it easier for search engines to index the publications so that they will return hits to them when people search for "cooperative auctions in robotics" or one of my other research topics. During your job search, your website should also contain pointers to material that is common to all your job applications (such as to your portfolios).

Managing Your Advisor

It is your job to manage your advisor. You want to talk to them early about their expectations of what it takes for you to earn your PhD degree (and take notes) and then talk to them every year again about your progress, what is left for you to do, and how you should best spend your time in the upcoming year. For example, when you join graduate school, you might want to discuss what kind of research output your advisor expects in terms of quality and quantity with respect to conference publications, journal publications, and systems. While you are at it, you might also want to discuss when your advisor will pay for your conference trips, how they determine (the order of) authorship on publications, when they expect you to be in your office or be reachable via phone or e-mail (for example, at night, on weekends, or during your vacations – the customs in academia vary a lot from research group to research group and might be very different from what you are used to), how quickly they want you to react to their requests, and how your vacations will work. Because you want to develop your own research agenda (with the help of your advisor), you should listen to them (after all, their job is to provide you with advice), but you do not always have to follow their advice[15] — although they will certainly influence your thinking via discussions. You need to learn how to disagree and argue with them in a convincing way. This skill is also important in writing papers for publication and your dissertation, defending your dissertation (where, as the name implies, you defend the importance of your hypothesis and the correctness of your arguments that prove or disprove it), and life after graduation. For example, if you think that your research is not yet ready for publication but your advisor feels pressure to publish it anyway (say, due to funding considerations), you might want to consult with a third party about the right thing to do rather than automatically giving in to your advisor. You might also want to find additional mentors, which could be more-senior members of your research group or researchers outside

of it. You should not be afraid to switch your advisor if your advisor does not give you enough freedom to pursue your own research agenda, does not have enough time for you, or is not willing to invest their time in you in other ways to make you successful (for example, by introducing you to other researchers or nominating you for awards). Having two advisors can be interesting if you work on interdisciplinary research and might mean more financial support, but also means that you might have to make two other researchers happy in addition to yourself.

Other Skills

We have discussed the most important skills in graduate school above. There are other skills — for example, being able to write grant applications and manage a team — but I consider them to be less important than the other ones because you will be able to acquire them later when you need them (although it has been pointed out to me that learning to negotiate is important before being on the job market because you will need to negotiate a good hiring package).

Conclusions

A PhD student stated: "Graduate school will be the best years of your life — in terms of the freedom to truly pursue anything you like and enjoy!" So, work hard, but also do not forget to have a life outside of graduate school as a counterbalance to the sometimes-stressful life in graduate school, such as during times when you feel that you are not making enough progress on your research! Don't stress yourself out. Have fun!

Acknowledgments

Many thanks to the following researchers who contributed comments to this publication: Liron Cohen, Jalal Kazemi, Stephanie Kemna, David Kempe, T.K. Satish Kumar, Ramesh Manuvinakurike, Alex Nash, Abdul Qadeer, and William Yeoh. Many thanks also to Rao Kambhampati and Wheeler Ruml for providing inspirations via the talks they gave in doctoral consortia before me, and to Wheeler Ruml for allowing me to use of the questions from his talk. Finally, many thanks to all researchers who provided me with advice over the years that I have incorporated into this publication and to the organizers of the IJCAI-19 doctoral consortium for their invitation!

Notes

1. This is the second habit in *The 7 Habits of Highly Effective People: Powerful Lessons in Personal Change* by S. Covey, New York, NY: Free Press, 2004.

2. I will also maintain an online version at idm-lab.org/advice.

3. I received inspiration and took the questions for the Understanding Potential Employers section of this article from AI researcher Wheeler Ruml's talks at the doctoral consortia of the International Conference on Automated Planning and Scheduling in 2008, 2010, and 2013. After graduation, Ruml worked first at an industrial research laboratory and then (and now) as a professor at a research university. You should look at his presentation Job Hunting in Industry and Academia (www.cs.unh.edu/~ruml/papers/job-hunting.pdf).

4. From A to Z of Computer Scientists by H. Henderson, *Facts on File* 2003. You can find Alan Newell's talk Desires and Diversions: 1991: Carnegie Mellon University online (www.youtube.com/watch?v=_sD42h9d1pk). After watching the video for a few minutes, I finally realized that I had been in the audience when he gave this talk at Carnegie Mellon University in 1991!

5. From *Models of My Life* by H. Simon, Cambridge, MA: The MIT Press, 1996.

6. For more information, see Progress on Agent Coordination with Cooperative Auctions by S. Koenig, P. Keskinocak, and C. Tovey, *Proceedings of the 24th Association for the Advancement of Artificial Intelligence* (*AAAI*) *Conference on Artificial Intelligence*, Menlo Park, CA: AAAI Press, 2010.

7. For more information, see Making Good Decisions Quickly by S. Koenig, *Institute of Electrical and Electronics Engineers* (*IEEE*) *Intelligent Informatics Bulletin* 13(1): 14–20, 2012 and Artificial Intelligence in 2027 by M. Gini, N. Agmon, F. Giunchiglia, S. Koenig, and K. Leyton-Brown, *AI Matters* 4(1), 10–20, 2018.

8. See *Artificial Intelligence — A Modern Approach* by S. Russell and P. Norvig, Upper Saddle River, NJ: Pearson/Prentice Hall, third edition, 2009.

9. You can find insights into the process of research in Uri Alon's talk Why Truly Innovative Science Demands a Leap Into the Unknown. (www.youtube.com/watch?v=F1U26PLiXjM&feature=youtu.be).

10. From Ethical Considerations in Artificial Intelligence Courses by E. Burton, J. Goldsmith, S. Koenig, B. Kuipers, N. Mattei, and T. Walsh, *AI Magazine* 38(2), 22–34, 2017.

11. From *Alan Newell* by H. Simon, Biographical Memoirs v.71, Washington, DC: National Academies Press, 1997.

12. I received inspiration for the Writing Publications and Dissertations and Giving Talks sections from AI researcher Rao Kambhampati's talk at the doctoral consortium of the International Joint Conference on Artificial Intelligence in 2013. You should look at his posted presentation, Wittgenstein's Papers and Faraday's Talks: Maxims for a Milk-fed Researcher (rakaposhi.eas.asu.edu/ijcai-dc-talk.html).

13. See also Dissertation Advice by Olin Shivers (www.ccs.neu.edu/home/shivers/diss-advice.html).

14. From Speeding Up Distributed Constraint Optimization Search Algorithms by W. Yeoh, PhD thesis, Department of Computer Science, University of Southern California, Los Angeles (California), 2010.

15. Of course, if your advisor provides you with research funding, then they have more influence on this part of your research, and you might want to look for synergies between your project work and your dissertation research — but your dissertation research is still yours!

Sven Koenig is a professor in computer science at the University of Southern California. Most of his research centers around techniques for decision-making (planning and learning) that enable single situated agents (such as robots or decision-support systems) and teams of agents to act intelligently in their environments and exhibit goal-directed behavior in real-time, even if they have only incomplete knowledge of their environments, imperfect abilities to manipulate them, limited or noisy perception, or insufficient reasoning speeds. Additional information can be found on his webpages at idm-lab.org.