

Ten Simple Rules for Starting a Research Group

Yolanda Gil

Information Sciences Institute and Department of Computer Science
University of Southern California
gil@isi.edu

October 27, 2014

One of my favorite discoveries in the last few years is a section of *PLOS Computational Biology* called “Ten Simple Rules” [PLOS Comp Bio 2014]. There are some real gems there, from “Ten Simple Rules to Combine Teaching and Research” [Vicens and Bourne 2009] to “Ten Simple Rules for Reviewers” [Bourne and Komgreen 2006], and my personal favorite “Ten Simple Rules for Better Figures” [Rougier et al 2014]. Although the articles are often specific to computational biology, many are written for a more general audience. I even humbly contributed to one as “Ten Simple Rules for the Care and Feeding of Scientific Data” [Goodman et al 2014]. The series is incredibly successful. The articles have been compiled as a monograph that has even been translated to Chinese.

Phil Bourne, a great model to follow in terms of community building and open scientific research, started the series as the editor of *PLOS Computational Biology* with “Ten Simple Rules for Getting Published”. He told me he wanted the series to represent concise and practical advice, with the name inspired by a TV show called “Eight Simple Rules”. For some reason, ten seems like a more scientific number and it has worked out beautifully for the series.

The initial article in the PLOS Computational Biology series was inspired by advice that resulted from a meeting with students, and so is this one. I recently participated in a Computing Research Association (CRA) meeting for Computing Innovation (CI) Fellows where I was asked to participate in a panel on “Starting Up a Research Group.” At this and similar panels in the past, I have found that students are always very appreciative of career advice. Many AI researchers already provide a wealth of useful advice and pointers on their web sites; perhaps having a series of articles on focused topics would be a more effective way to disseminate that advice.

I personally appreciate all the advice from the “Ten Simple Rules” articles. Scientific research methods evolve, and we must be able to distill and share best practices as they emerge. There are active discussions in the community on open access publication, alternative approaches to conference reviewing, improved

reproducibility by publishing papers that include data and code, and broader views of what is considered a scientific contribution. Students also have much to teach us about using social media for scientific organizations, committing code, and tweeting during conferences (wait, there already is “Ten Rules of Live Tweeting at Scientific Conferences” [Ekins and Perlstein 2014] for that!).

For all these reasons, I thought we could use an analogous “Ten Simple Rules” series catering to AI researchers in the new ACM SIGAI *AI Matters* newsletter. Phil was very pleased to hear this, and noted that for the upcoming anniversary of the collection there would be “Ten Simple Rules for Writing a PLOS Ten Simple Rules Article” [Dashnow et al 2014].

I encourage you all to read the existing “Ten Simple Rules” collection and to consider submitting your own “Ten Simple Rules” articles for future issues of this newsletter. Or ask your favorite mentors or students to submit one.

Rule 1. Be in the Right Place

There is nothing worse than working very hard but finding that your work environment is not helping you succeed. Imagine finding yourself in a place where there are too many unwanted obligations and demands on your time, or you do not have the collaborators that you need, or you do not feel valued. You want to avoid this at all cost.

When looking for a position, you should think about how that institution is going to help you further your career goals. It should provide resources that will complement your skills and goals. Find out how younger researchers are supported, for example through senior mentors or startup packages that enable them to fund students and travel. Your institution should be a place that helps you learn and grow, with a culture that suits your style. If you are in the right place for you, you will be much more motivated and productive.

When I graduated, I thought very hard about what kind of position suited me best. I enjoy being very independent and like to interact with students. But I did not want to be spending a big chunk of my time teaching every semester. I




Figure 1. Be in the right place. Your institution should be a place that helps you learn and grow, with a culture that suits your style. If you are in the right place for you, you will be much more motivated and productive.

decided to take a job at a university research lab, and this allowed me to pursue my own research agenda full time while being in an academic environment. I have been here for more than twenty years.

Rule 2. Formulate an Exciting Vision and Clear Research Plan

Many people will ask you about what you are working on. It is easy to fall back into your thesis and talk about how you are continuing that work. But you may want to do more and show that you have ideas independent of your advisor. More than that, you want to show that you have big ideas so you can foster a reputation as an interesting colleague and as a thought leader in your field.

The single most important key to successfully starting a research group is to formulate an exciting vision and a clear research plan to accomplish it. This takes a lot of thought, particularly if the research is innovative and transformative or if you are departing from the area in which you did your thesis work. The time you invest in this will really pay off. You will be able to tell others clearly what you are doing: your prospective students, potential funders, reviewers, etc. You will have a much easier time growing your research group (see Rule #5 and Rule #6).

Learn to present your research. Whether you are giving a talk or writing, you must convey your excitement about making progress on important problems.

When I started my career, I found it very useful to look at the research plans of more senior people. Most often research plans only exist in people's heads, but when people apply for a new position (or a promotion/tenure), they will write them down in a coherent and exciting way. These days you can often find them on people's Web sites, but you can also directly ask researchers whose work you admire.

Rule 3. Define Short- and Mid-Term Goals

If you are not driven by the things you want to accomplish, your time will be taken over by other things. You will get many offers to be involved with committees, reviewing, and student events. These are important, but you must keep your focus.

You should think about your own goals, short term (six months out) and mid-term (six years out) and long term (sixty years out). Your goals should include everything from doing research to gathering resources such as students and funding to disseminating your findings. Adjust these goals if good opportunities arise, but always have a set of goals that are your own. Otherwise, you will get distracted and accept too many impositions on your time.

I plan in six month periods, making very crisp six month goals, such as submitting papers to specific conference deadlines, working on particular research questions,

and traveling to attend selected conferences or to give talks at specific venues. I make sure that the goals are feasible within six months. I make sure everything else stays in the drawer of future to-dos. Every six months, I open that drawer and also revisit my long-term goals, consider my priorities, and plan for the following six month period.

Rule 4. Multiply Your Own Productivity

Many of the things we do take a long time to prepare, whether making a new presentation, filing a project report, or adding content to our web sites.

Make everything you do count twice, so that it is useful for more than one purpose. If you take a trip to a conference, think of a place nearby where you could give a talk about your research. If you write a paper based on the work for a project, it could be reused as a project report or for the project web site. If you put effort in preparing a new talk, you might as well write a paper based on that talk.

This article is an example of what I mean: It is based on a presentation I gave recently at a CRA meeting and will give again soon to Ph.D. students at my university. Since I had already made the effort to organize the presentation, writing this article was easy.

Rule 5. Find Special Funding Opportunities

Funding allows you to hire students and grow your research group. But at the beginning of your research career, getting funding is challenging. Writing good proposals takes a lot of skill, and getting them funded involves careful strategy and a fair amount of luck. Of course, it all starts with having exciting ideas (see Rule #2).

Funding agencies in general make supporting new investigators a priority. Many offer CAREER programs for young faculty. NSF has a new Computer and Information Science and Engineering (CISE) Research Initiation Initiative (CRII) program for early career researchers. NIH has special career development awards




Figure 2. Define short- and mid-term goals. Your goals should include everything from doing research to gathering resources such as students and funding to disseminating your findings. Otherwise, you will get distracted and accept too many impositions on your time.

for post-graduate researchers, particularly in its Big Data to Knowledge (BD2K) program where they seek to attract computer scientists. If you do not know of such programs in your area, ask your colleagues or ask program managers in funding agencies in your research area. Your institution is also likely to have programs to help researchers find appropriate funding mechanisms and learn to write grants.

One way to get a funding stream started is to submit proposals jointly with colleagues who have more experience. You will learn a lot from them, and in the process you may get some funding to get your research group started. If you have a hard time formulating joint proposals, at least seek others' feedback on your own proposals.

Another way is to seek internal funding at your institution. Many universities have small internal grants for faculty, and they may be more likely to give them to young faculty like you since they want to help you be successful.

Whatever your strategy, do not be despondent if your proposals are rejected. Just keep improving them and resubmitting better and better versions. Your persistence will eventually pay off.

Rule 6. Seek Collaborators and Grow Your Team

If you work by yourself, you will be able to get less done than if you have a team of students and colleagues working with you. But when you are starting out, you often do not have funding to form a team.

There are a few ways to grow your research group even without funding. If you have articulated a promising research agenda (see Rule #2), you should be able to attract collaborators who will be excited to work on it. You may attract colleagues to work with you, perhaps even spending a sabbatical or a long visit at your institution. It may also be possible for you to get students even if you do not have funding. You can work with graduate students who are supported by fellowships. Undergraduate students often make themselves available for unpaid summer internships. Students from foreign universities sometimes have grants to do research abroad. With collaborators, you can multiply your productivity.

Over the years I have had many collaborators that I have not supported with funding. I have had several visiting professors on sabbatical, several graduate students with scholarships, undergraduates interested in learning to do research, and even high-school students wanting to get some experience. Personally, I find it tremendously rewarding to work with undergraduates and high-school students; if they are given the right preparation and the right task they are very productive and creative.

Rule 7. Learn To Work with Others

People have their own minds, they will pursue things that matter to them, and they do not always do things just because you ask them to. This may lead to frustration on your side as you try to build a team of collaborators that is productive and enjoyable. Keeping people happy and motivated is tricky, and it is not a skill that you necessarily learn in graduate school where most people work independently on their own thesis research. As soon as you graduate, you will most likely have to work closely with other people as collaborators and students, and occasionally with colleagues in faculty or conference committees.

To work effectively with collaborators, you have to learn to work with other people. Read books about this, observe your best collaborators, or simply use common sense. Here are a few tips to get you started. Do not approach people by asking them to do things because you need them done. Instead think about ways in which they will benefit from the work and emphasize that. Do not expect people to get things done unless you set a deadline or at least a timeline, as people are often too busy and will end up giving priority to other things that have more urgency (and of course, you should send a reminder as the deadline approaches). If you are organizing a group and each person has a task, lead by example and get it done, then make sure everyone knows when someone gets theirs done so they feel the peer pressure. Finally, only bother people when it is something important because if you bother them with big and small things they might clump them all together and defer.

You should also learn to manage people. If you work in a company or research lab you will have many chances to put your management skills to use. These skills are useful even if you are in academia, because there is a good chance that you will have a group of students to manage. A good manager listens and observes and has a good model of how to motivate each person in his or her team. A good manager also understands the weak points of each person and when possible helps them learn to improve. We all know that nobody is perfect, but it may take time for us to realize that we often judge people as if they must be so. You may get frustrated because other people are not getting things done exactly the way you would do them, but if you want to work with other people you have to accept that and be okay with it. Your job is to carve out work for each person in your team that they will find interesting and well aligned with their motivations and skills.

I learn a lot at social gatherings about how to manage people. I ask people about the hardest situation they handled, or what they do when someone does not work well within their team. Another great source for me has been the advice for managing volunteer groups in open source software communities or standards bodies. I never liked management and I do not think I was cut out to be a manager, but I have become much better at it over time. You can too.

Rule 8. Add Good Mentors to Your Team

Sooner or later you will face a challenging situation where you will not be sure what to do. Perhaps you are not sure what might be a more productive research avenue. Perhaps you are being asked to review for many things and you are not sure how to manage so many commitments. Perhaps one of your collaborators did not include you in a paper and you thought you had contributed enough to be included. Perhaps one of your best students seems to be disengaged and you are worried that they will drop out. Whatever it is, chances are that you are not the first person to experience this challenge. So seek advice from those who came before you.

Think of the people that you look up to and make them your mentors. This does not have to be an official arrangement; just make a point to get together and talk to them regularly. They do not have to be at your institution; they can be people that you see at conferences. Set up a lunch with them and ask about everything that is on your mind. You will be pleasantly surprised by what good suggestions they have, or simply how much they know about how others handle the same issues. Sometimes they will give you a better perspective that will help you handle whatever looks overwhelming. You cannot learn to do research from a book, so mentors are your best way to gather knowledge about how to best do your job.

I have always found that people are eager to give advice and help if you ask them. Very few people will volunteer advice if you do not ask first. And I have always learned something useful that helped me with whatever matter I was concerned about. The biggest bonus is when I get unsolicited advice about matters that I was not even thinking about.

Rule 9. Have a Broad View of “Team”

Sometimes you think about your team as your students and your research group. But think about how many colleagues could be on your team as collaborators or simply as supporters. Imagine if you could rely on a personal network of supporters that could give you comments on paper drafts, write letters of support when you need them, and generally be a sounding board for your research.

Create a network of people who know and appreciate your work. Think about who would be in your “dream team” and initiate a conversation. Talk about why you like their research and how your work is related to theirs. Make a point to meet with them at conferences. Read their papers regularly and send them your thoughts. Meet their students whenever you can and ask them to apply for a summer internship in your group. If appropriate, invite them for a visit and to give a seminar. You may get a mentor or two in the process (see Rule #8).

I was always shy about imposing on busy people’s time. Then I realized that busy people are, like everyone else, eager to hear good ideas and meet the best people in

their field. Do not hesitate to approach people whom you think should be aware of your research, no matter how busy they appear.

Rule 10. Keep It All Sane

Being a researcher means setting up your own work plan and pace, and it is so much fun that we often overcommit and overbook our time to the point of being insanely busy. We have a profession that is highly prone to “workaholism”. When you start your career, you are particularly at risk for it. You are young and believe you can endure anything: the late nights, the deadlines, the unexpected problems: you can handle all that. At some point, sometimes sooner and sometimes later, you will realize this is not a good way to live. Perhaps you are not spending enough time with your family. Or you act like a zombie after working through many late nights. Or you are too exhausted to generate your best ideas.

In the words of one of my mentors, our work is almost always “self-imposed” and we just have to control ourselves. You can take charge and control the madness. Watch yourself, and do not commit to doing things unless you can reasonably make the time to do them. Another one of my mentors invented the “not-to-do list”, where you put in a very visible place (as visible as your to-do list) that enumerates all of the things that you declined to do. When you feel overwhelmed, look at your not-to-do list and you will be suitably reinforcing the good decision that you made to not overcommit. Listen to your body. Instead of always letting your mind tell your body what it must endure, let your body speak up sometimes about whatever rest it needs. Vacation is never overrated, no matter how much you love your work.

Avoid a reputation of agreeing to do things and then dropping the ball so that others have to step in at the last minute and pick up the pieces. Avoid regrets because you were not able to go to the game with your kids. Controlling the madness should be your foremost priority. Do get a life.

Whatever time I spend away from work makes me come back to it with more zest for the excitement of doing research.



Figure 3. Keep It All Sane. In the words of one of my mentors, our work is almost always “self-imposed” and we just have to control ourselves. You can take charge and control the madness. Watch yourself, do not commit to doing things unless you can reasonably make the time to do them.

Acknowledgments

I would like to thank my colleagues for their comments on earlier versions of this paper, in particular Jonathan Gordon, Jeon-Hyung Kang, Kevin Knight, and Jonathan May, and Ashish Vaswani.

References

- [Bourne 2005] “Ten Simple Rules for Getting Published.” Philip E Bourne. *PLOS Computational Biology*, 1(5), 2005.
- [Bourne and Korngreen 2006] “Ten Simple Rules for Reviewers.” Philip E Bourne, Alon Korngreen. *PLOS Computational Biology*, 2(9), 2006.
- [Dashnow et al 2014] “Ten Simple Rules for Writing a PLOS Ten Simple Rules Article.” Harriet Dashnow, Andrew Lonsdale, Philip E. Bourne. *PLOS Computational Biology*, 10(10), 2014.
- [Lonsdale et al 2014] “Ten Simple Rules for Writing a PLOS Ten Simple Rules Article.” Harriet Dashnow, Andrew Lonsdale, and Philip E. Bourne. *PLOS Computational Biology*, 10(10), 2014.
- [Ekins and Perlstein 2014] “Ten Simple Rules of Live Tweeting at Scientific Conferences.” Sean Ekins and Ethan O. Perlstein. *PLOS Computational Biology*, 10(8), 2014.
- [Goodman et al 2014] “Ten Simple Rules for the Care and Feeding of Scientific Data.” Alyssa Goodman, Alberto Pepe, Alexander W. Blocker, Christine L. Borgman, Kyle Cranmer, Merce Crosas, Rosanne Di Stefano, Yolanda Gil, Paul Groth, Margaret Hedstrom, David W. Hogg, Vinay Kashyap, Ashish Mahabal, Aneta Siemiginowska, and Aleksandra Slavkovic. *PLOS Computational Biology*, 10(4), 2014.
- [Pavelin et al 2014] “Ten Simple Rules for Running Interactive Workshops.” Katrina Pavelin, Sangya Pundir, and Jennifer A. Cham. *PLOS Computational Biology*, 10(2), 2014.
- [PLOS Comp Bio 2014] Ten Simple Rules. *PLOS Computational Biology*. Available from <http://www.ploscollections.org/article/browse/issue/info%3Adoi%2F10.1371%2Fissue.pcol.v03.i01>.
- [Rougier et al 2014] “Ten Simple Rules for Better Figures.” Nicolas P. Rougier, Michael Droettboom, Philip E. Bourne. *PLOS Computational Biology*, 10(9), 2014.
- [Vicens and Bourne 2009] “Ten Simple Rules To Combine Teaching and Research.” Quentin Vicens, and Philip E. Bourne. *PLOS Computational Biology*, 5(4), 2009.