Guidelines for working in the Numerical Mathematics Group at KAUST

February 23, 2013

Welcome to the Numerical Mathematics group! This document is a set of "best practices" and helpful advice for working in the group. These guidelines are written primarily with graduate students in mind, but most of the contents are also useful for post-docs and research scientists in the group.

1 Meetings

1.1 Weekly group meetings

We meet as a group once each week over lunch. This meeting is meant to be a social time to discuss both “real life" and research informally. Usually there will also be a structured presentation by a group member about their research, or a discussion of a paper that we’ve all read, or discussion of important skills and tools. You’re expected to give a presentation or lead a discussion at least once each semester. This is your opportunity to develop your presentation skills in a friendly environment where you’ll receive only helpful criticism. When you’re not presenting, it’s your opportunity to ask lots of questions.

1.2 Weekly individual meetings

In addition to weekly group meetings, I usually hold weekly individual meetings with everyone in the group. Depending on your research focus, it may make more sense for you to meet weekly with a research scientist or post-doc in the group, and less frequently with me. In any case, these meetings usually last one half hour, but longer meetings can be arranged when necessary. This is your time and the meeting is really driven by you. Before the meeting, you should take a few minutes to prepare an **agenda**, listing all the items that you'd like to discuss. I will often have some to add. It's helpful if you e-mail me a copy of your planned agenda before the meeting.

The first purpose of the weekly meeting is for you to update me on your research progress. At the same time that you prepare your agenda, you should prepare a **short summary** (often a paragraph is sufficient, but you may wish to include plots, tables, etc.) describing what you have accomplished in the last week. This typically includes things like reading papers/books, writing, programming, and working through analysis of a problem. Sometimes this will be "I tried approaches X, Y, and Z but none of them worked"; that's fine. The summary isn't for me to read during the meeting, but as a tool to guide the conversation and a permanent record of what you've been working on that is sometimes useful later on.

The next purpose of the meeting is for you to ask me about points that you don't understand or approaches that aren't clear. You should try to come with at least one question every week. I will usually sketch an answer and provide you a reference with more detail. As you progress and become the expert in your own research, the tables will turn and I will start asking you to explain things to me.

The third main purpose of the weekly meeting is to plan your work for the next week. You should come to the meeting with a short **written plan** for the following week, which may be added to or modified during the meeting. You should always leave with a very clear idea of what you are going to do. If you don't have that, tell me so. I will often commit to do some things as well; if you notice that I haven't done something we agreed on, please remind me.

Finally, the weekly meeting is a convenient time to discuss any other necessary business, such as forms to be signed, travel requests, equipment, exam preparation etc. These things often come up during the week; unless they are urgent, it is preferred to discuss them all at our one-on-one meeting.

1.3 Seminars

Go to seminars. Not just the ones that seem directly related to your research; go to lots of seminars, including some in other disciplines that look interesting. You'll find out about new problems and new techniques, learn how (or how not) to give a good talk, and stretch your mind.

2 Research

While we usually focus on the direct results of research, there are many important skills that you will acquire indirectly as you begin doing research:

- Absorbing and navigating the 'big picture' of a field/subfield
- Learning to read and understand the literature
- Learning to write technical material clearly and precisely
- Learning to communicate verbally, by asking questions and by presenting research
- Learning to formulate and assess new research questions, methods, and results

At first the research itself is secondary to developing these skills, because without them your research will not have impact.

2.1 Finding a thesis topic

You should begin discussing your research interests and style with me from day one. Besides subject matter, the following aspects are also important: Are you interested more in analysis, applications, or implementation? Do you like proving theorems or writing code? Open-ended exploratory projects or focused questions with specific answers? Do you like working closely with others or having a project that you will advance independently?

For a MS thesis or directed research, I will provide you with a project and desired objectives.

For a PhD thesis, I will provide ideas and references for you to explore based on the intersection of your interests with my own. You should also do your own independent searching and reading on topics that you're interested in. Eventually (usually by the

end of your first year as a Ph.D. student) you should (in consultation with me) settle on something that looks like a promising thesis topic. An ideal thesis topic will include both clear, incremental advances that are almost certain to be worthwhile as well as some extension(s) that may or may not work but would be very valuable. Once the thesis topic is identified, you should prepare a thesis proposal. The proposal need only be a few pages long, identifying some clear objectives and also indicating what is *not* expected to be part of the thesis.

Inevitably, some objectives or methodologies of the thesis will change from those originally envisioned, usually due to complications in the research that cannot be seen at the outset. These new objectives and methods will be invented by you, as you become the leading expert in your thesis topic.

2.2 Planning

I have already mentioned the weekly plan that you should bring to our meetings. To keep your research on track, you are expected to also prepare a plan at the beginning of each semester. The plan should include your research objectives for the semester and a timeline for accomplishing them. We'll have a group meeting early in each semester devoted to discussion of research plans.

2.3 Lab Notebook

I strongly recommend that you keep a research notebook. This can take many forms, but it is probably best to use something electronic. It's a good idea to keep notes as you go along throughout the day, but at a minimum you should make an entry at the end of each day. Write down what you've done, especially making sure to include mention of any code you've writtten and how to use it. Write down answers you obtained or surprising things you observed. Write down what steps need to be taken next. The writing will help clarify ideas in your own mind and help you remember what you've done when you need to revisit it later on.

2.4 Reading

Learning to read and understand the literature related to your research is an essential skill, and one that will also help you develop your writing (see below). Some important

parts of it are described in an entry on my blog. You can also see some suggestions of journals to follow in this blog entry.

You can see what I'm reading by following my Mendeley feed. In fact, this is a good way to find new papers that I consider to be potentially important. I don't actually get time to read most of those papers in detail, so when you do read one of them, make a summary of your notes on it and give them to me. I'll be grateful.

You should be actively searching the literature on topics related to your research and sharing useful references you find with other members of the group who are involved in similar topics.

2.5 Writing

In mathematics and other related sciences, technical writing is done with laTeX. If you haven't learned to use laTeX, it should be one of your first priorities. You'll also want to learn how to use BibTeX to manage references. It's a good idea to use some kind of desktop software (like Mendeley) to facilitate management of your bibtex references.

You will find it wothwhile to invest time in developing excellent technical English writing skills. This is especially true if (as for most researchers) English is not your first language. I have several books on mathematical writing that you are welcome to borrow. There is also an office on campus that can assist you with your writing. Correct usage is essential, but you will also want to develop good style, which is a skill you will continue develop throughout your career. Some potentially useful resources in this regard are:

- `http://terrytao.wordpress.com/advice-on-writing-papers/`
- `http://www.stat.rice.edu/~riedi/HalmosWrite.rtf`
- `http://www.math.uga.edu/~azoff/courses/halmos.pdf`
- My personal favorite, despite the title's punctuation: `http://www.damtp.cam.ac.uk/user/ai/HowTo07.pdf`
- `http://www.ams.org/notices/200711/tx071101507p.pdf`
- `http://ems.calumet.purdue.edu/mcss/kevinlee/mathwriting/writingman.pdf`

You should start writing as soon as you start doing research. Most of your writing will not go into any publication, but putting your ideas down carefully in writing will help you discover gaps and show you what to think about next. It will also help you

to remember what you have already done, which is surprisingly difficult after a few months.

2.6 Speaking

Take advantage of all the opportunities you can find to give presentations. These include our weekly group meetings and student events on campus (SIAM chapter seminar series, annual student research symposium, etc.) As soon as you have got even a small interesting result you can begin presenting on these kinds of occasions. If you're like me, it will take you several tries to get beyond natural nervousness and be able to focus well on what you're presenting.

After you've given presentations to these friendly audiences, and once you've got some substantial research under your belt, you should think about giving a presentation at an international conference. I may suggest such an opportunity to you, but you need not wait for me to do so. Be aware of upcoming conferences that might be appropriate places to present your work.

Before giving a presentation (with the exception of informal presentations to other students or at our research group), you should show your slides to me and to any co-authors of the work you will speak about.

2.6.1 Answering questions

The question and answer period following your talk is often even more important than the talk itself in terms of the impression you leave with an audience.

When you're asked a question, carefully rephrase what you understand the question to be before answering. It's always okay to admit that you don't know the answer to a particular question, and it's much better than trying to fake it.

It is natural, but wrong, to react to questions as if they were an attack on your work. Never get defensive or antagonistic with a questioner. If you start to feel this way, politely suggest that the discussion be continued later. Arguing with someone in a public forum (even if you're right) is a sure way to leave a poor impression.

2.6.2 Improving your style

Once you've given a few technical presentations, you'll be able to present more clearly and with less nerves. Then it's time to start improving your style. Giving engaging

presentations of mathematical material is challenging and many mathematicians are boring to listen to. The most important first steps are:

1. Give an appropriate amount of technical detail. This depends on the kind of talk you're giving, but almost always means less detail than you would naturally wish to include. Listeners cannot absorb very complex material during a presentation.
2. Don't write your talk on your slides, and don't read your slides to your audience. Do create good figures that convey your ideas without many words.
3. Minimize the amount of mathematical notation, and introduce it clearly.

Next, you can work on becoming a really polished presenter. I recommend the following resources:

- `http://research.microsoft.com/en-us/um/people/simonpj/Papers/giving-a-talk/giving-a-talk-html.html`
- *Even a Geek can Speak*. The E-book is available through Amazon for just $10. It's aimed at businesspeople, but almost all the advice carries over to research talks.

2.7 Publishing

Journal publication is an essential part of science, since it is the recognized medium for disseminating research in detail. It is also an essential part of an academic career, since it is the primary metric used in evaluating researchers. Never forget that the use of publishing as a metric is incidental, whereas its real purpose is the communication of ideas.

Early in your thesis research, you should not be concerned with publishing at all; trying to publish a journal article too soon can even be harmful because you may end up publishing an incomplete result that you are later able to improve. Depending on the subfield you are in, publishing in a conference proceedings early on may make sense. As you move closer to graduating, publishing should be a priority if you are looking for an academic or government post-doc, since employers will expect you to have some publications. As always, quality is more important than quantity.

Broad scientific norms, as well as KAUST's own policy on research conduct, stipulate that all authors listed on a paper should have made substantial contributions to both the research and the writing, and that all persons who have made such contributions to both should be listed as authors. In applied mathematics, it is typical to list authors in order of their contribution.

2.8 Programming

At the moment, most programming in the group is done in Python, Fortran, or MATLAB (roughly in decreasing order of importance). You'll need to learn one or more of these if you don't know them already. Again, don't be afraid to ask for help from your colleagues.

If a significant portion of your work will involve programming, it is worthwhile to invest time in learning how to do it well. This time will repay itself many times over. During your first semester in the group you should read one or both of the following:

- Code Complete (Steve McConnell)
- The Pragmatic Programmer: From Journeyman to Master (Dave Thomas and Andy Hunt)

I have copies of both that you may borrow. You may also find it worthwhile to purchase a copy.

2.9 The Big Picture

You should strive to understand the larger context of the problem you are working on within your subfield and within the scientific realm. Doing so will help you accurately estimate its importance and recognize other related significant problems. Distinguishing important problems is one of the most important skills for any scientist. Read Richard Hamming's essay *You and your research*, and be sure to devote time regularly to "great thoughts".

3 Other Skills and Tools

3.1 Unix

You should use a Unix-based operating system; I recommend either Mac OS X (which is for now the most widely used in our group) or Ubuntu (note that you can run Ubuntu on a Mac using a virtual machine). If you're not already familiar with the Unix shell, you'll want to learn to use it effectively. Find another member of the group (or other colleague) who can answer your questions, and don't be hesitant to ask. Effective use of the shell can make you much more productive in many day-to-day tasks.

3.2 Text editors (Vim or Emacs)

You will spend a substantial amount of your research time writing code, papers, and reports, so you will want to make use of a good text editor, like Emacs or Vim. Editors like these dramatically reduce the amount of typing needed in order to cut, paste, and rearrange text within a file. I use Vim, but Emacs is equally useful (on the Mac, MacVim is especially nice). I've also heard good things about TextMate. For either, you can find "cheat sheets" that will help you quickly learn the most useful commands. If you find a better text editor than these, let me know.

3.3 Version Control (Git, Mercurial)

Version control software is an essential tool that we use both for developing code and for writing papers. It allows different members of the group to work simultaneously on the same code/paper and it also simplifies tracking changes. It also saves you from making ten copies of the same directory without any indication of which is the "definitive" version. Thus you may find version control useful even for your own projects that don't involve anyone else. We primarily use Git. There are several good Git tutorials on the web. For public repositories, we typically rely on Github for hosting; for private repositories, we use Bitbucket.

4 Miscellany

4.1 Webpages

The group website is `http://numerics.kaust.edu.sa`. There is space for each group member to have a page. You are responsible for preparing your own, using one of the existing pages as a template. It's a good idea to post things like slides from talks that you give on your webpage. In any case, you should at least send me a minimal amount of information to place on the "People" page, including a photo and your research interests.

I post announcements about significant happenings within the group (publications, workshops, etc.) on the main page. If you know of something that should be announced there, let me know.

If you would like to volunteer to take over management of the group website for me, let me know.

Some great advice on doing mathematics can be found at http://press.princeton.edu/chapters/gowers/gowers_VIII_6.pdf

4.2 Conduct

Although academic culture is often seen as hierarchical, you should view everyone in the group as your colleague and potential collaborator. You should treat everyone with whom you interact professionally with courtesy and respect, regardless of their job, rank or experience. Be especially careful about the tone of e-mails that you send, because they are more easily misinterpreted than face-to-face conversations.

4.3 Backups

This not optional: use a backup drive. Get a big drive (you can be reimbursed for the cost) and turn on Time Machine. Someday your hard drive will fail, and when it does you (and I) will be very glad that you did this.

5 Final thoughts

Scientific research is a highly creative process and requires self-motivation. Your goal should be to to find and pursue research questions that make you excited about getting up in the morning. My role is to guide you in the process of discovering and answering those questions. Some temporary slumps are normal in the course of a long research project, but if for an extended period you are not excited about what you're doing, come talk to me.