

Course notes for MIT EECS 6.819/6.869 Spring 2022

Antonio Torralba, Phillip Isola and William Freeman

April 25, 2022

Contents

41 How to do research	5
41.1 The Transition From Classes to Research	5
41.2 Toy Models	6
41.3 Strategies for Success	7
41.4 Collaboration and Authorship	8
41.5 Closing Comments	8
42 How to Write Papers	9
42.1 Motivation	9
42.2 Organization	9
42.2.1 Introduction	10
42.2.2 Main Idea	10
42.2.3 Experimental Results	11
42.2.4 How to End the Paper	11
42.3 General Writing Tips	11
42.3.1 Use Fewer Words	11
42.3.2 Title, Figure Captions, and Equations	13
42.3.3 Tone	15
42.3.4 Author List	15
42.3.5 Avoiding Rejection	15
43 How to Give Talks	17
43.1 Motivation	17
43.2 Prepare	17
43.3 Nervousness	17
43.4 Very Short Talks (2 - 10 minutes)	18
43.5 Longer Talks (15 - 60 minutes)	18
43.5.1 Your Audience	18
43.5.2 Ways to Engage the Audience	18
43.5.3 How to End the Talk	20

Chapter 41

How to do research

41.1 The Transition From Classes to Research

The jump from problem sets to research can be hard. We sometimes see students who ace their classes struggle with their research. In little bites, here is what we think is important for succeeding in research as a graduate student. This chapter is written as advice from research advisor to a graduate student, so we write as "I" in the bullet points below.

- The first advice can go on a bumper sticker: "Slow down to speed up". In classes, the world is rigged. There's a simple correct answer and the problem is structured to let you come to that answer. You get feedback with the correct answer within a day after you submit anything.

Research is different. No one tells you the right answer, and we may not know if there is a right answer. We don't know if something doesn't work because there's a silly mistake in the program or because a broad set of assumptions is flawed.

How do you deal with that? Take things slowly. Verify your assumptions. Understand the thing, whatever it is—the program, the algorithm, or the proof. As you do experiments, only change one thing at a time, so you know what the outcome of the experiment means.

It may feel like you're going slowly, but you'll be making much more progress than if you flail around, trying different things, but not understanding what's going on.



Figure 41.1: Research advice for a bumper sticker.

- Please don't tell me "it doesn't work". Of course, it doesn't work. If there's a single mistake in the chain, the whole thing won't work, and how could you possibly go through all those steps without making a mistake somewhere? What I want to hear

instead is something like, “I’ve narrowed down the problem to step B. Until step A, you can see that it works, because you put in X and you get Y out, as we expect. You can see how it fails here at B. I’ve ruled out W and Z as the cause.”

- “This sounds like hard work.” Yes. It’s no longer about being smart. By now, everyone around you is smart. In graduate school, it’s the *hard workers* who pull ahead. This happens in sports, too. You always read stories about how hard the great players work, being the first ones out to practice, the last ones to leave, etc.

(My co-author, who generally works harder than I do, tells me I should add comments here about the importance of taking time off from work and maintaining a good life/work balance. That is true, and you should work at a pace you can sustain, and you should refresh yourself by making time for the things that matter more than work, such as relationships, service, and relaxation. Perhaps the point is to protect your time so you will have enough time to do the research well.)

- “How do I get myself to work hard enough to do research well?” It all plays out if you love what you’re doing. You become good at it because you spend time at it and you do that because you enjoy it. So pick something to work on that you can love. If you’re not the type who falls in love with a problem, then just know that working hard is what you have to do to succeed at research.
- I have to note that the above isn’t completely true. Beyond working hard, there’s also *steering*. We’re like boats. We need motors—that’s the working hard part. But we also need a rudder for steering—that’s stepping back periodically to make sure we’re working on the right thing. On the topic of steering, I find time management books to be very helpful. They teach you how to spend your time solving the right problems.

41.2 Toy Models

- There’s a concept I want a simple phrase for, and maybe you can help me think up a good name. It’s the simplest toy model that captures the main idea. TSTMTCTMI? Anyway, simple toy models always help me. With a good one, you can build up intuition about what matters, which is a big advantage in research.

Here’s an example. The color constancy problem is to estimate surface reflectance colors when we only get to observe the wavelength-by-wavelength product of the each surface reflectance spectrum and the unknown illuminant spectrum. A toy model for that problem is to try to estimate the scalars a and b from only observing their product, $y = ab$. There’s a surprising richness even to this simple problem, and thinking about it allows you to think through loss functions and other aspects of Bayesian decision theory. I co-authored a paper that discusses $y = ab$ for much of the manuscript. Another toy model: as a proxy for complicated shaded surfaces, a single bump. You get the idea. Having the intuitions from working with toy problems gives you a big




Figure 41.2: A toy model for the color constancy problem: $y = ab$

advantage in the research, because you can figure out what will work by thinking it through with your toy model.

41.3 Strategies for Success

- A parable, as told by my friend Yair Weiss: There is a weak and a strong graduate student. They are both asked by their advisor to try a particular approach to solving a research problem.

The weak student does exactly what the advisor has asked. But the advisor's solution fails, and the student reports that failure.

The strong student starts doing what the advisor has asked, sees that it doesn't work, looks around within some epsilon ball of the original proposal to find what does work, and reports that solution.



Figure 41.3: The parable of the two students

- Sometimes it's useful to think that everyone else is an idiot. This lets you do things that no one else is doing. It's best not to be too vocal about that. You can say something like "Oh, I just thought I'd try out this direction".
- It's also sometimes useful to remember that many smart people have worked on this and related problems and written their thoughts and results down in papers. Don't be caught flat-footed with a large body of closely related literature that you aren't familiar with.
- Here's how a business school might talk about your research. You have a brand: you. There are many impressions you want to build up about your brand: that person always does great work, they have good ideas, they give great talks, they write wonderful software. Promote your brand. Build up a great reputation for yourself by consistently behaving as you'd like to be thought of.



Figure 41.4: Nurture your research brand.

- Cultivate your strengths and play to those strengths. Some possible strengths: being broad; creative; a great implementer; great at doing theory.
- Please don't report to me, "This instance doesn't work". Why doesn't it work? Why should it work? Is there a simpler case we can make work? Do you think it's a general

issue that affects all problems of this category? Can you think of what's not working? Can you contort things to make an example that does work? At least, can you make it fail worse, so we understand some aspects of the system?

- I love to hear about progress when I meet with students, but note that I have a very general notion of progress. Progress can include: "I've shown why this doesn't work", "I've simplified the task to get it to start working.", or "I spent the whole time reading because I know I have to understand this before I can make any progress."
- Please don't hide from me. Let's talk. I like it when you track me down and insist that we talk, for example, if I've been traveling.

41.4 Collaboration and Authorship

- On Collaboration: Science is generally a team sport. What matters in a collaboration is whether the work is insightful, foundational or impactful. It doesn't matter that many people were involved. It is much better to be one of 10 contributors to a great paper than to be the sole author on an unimportant paper. The relationships formed through collaborations can become friendships that last over your career.
- On Authorship: What names should be on a paper? Everyone who contributed to the paper. Contributions can be through experimentation, contributing a seminal idea that the paper builds on, or helping with the writing. Sometimes, an authorship contribution can include trying out a research direction that didn't work,

If I'm on the fence about whether someone contributed enough to be an author, I usually ask the person themselves, and go with whichever outcome they prefer.

41.5 Closing Comments

- For a presentation to the visiting admitted MIT Computer Science graduate students, I emailed all CSAIL researchers and faculty members, "Please send me what you think is the most important quality for success in graduate school" I compiled their responses (along with photos of the responders) into slides that are available online:

<http://people.csail.mit.edu/billf/talks/10minFreeman2013.pdf>

I think it's a lot of good advice about research.

- One final note about doing research: I hope you love it. I certainly do. The research community is a community of people who are passionate about what they do, and we welcome you to it! ¹

¹Photo credits:

1. <http://bethpartin.com/denver-photos-vw-bus-on-capitol-hill>
2. D. H. Brainard and W. T. Freeman, Bayesian Color Constancy, *Journal of the Optical Society of America, A*, 14(7), pp. 1393-1411, July, 1997
3. http://www.123rf.com/photo_2851187_the-two-students-with-the-book.html
4. <http://hms-somerset-co.blogspot.com/2012/12/apple-logo.html>

Chapter 42

How to Write Papers

42.1 Motivation

An important part of research is communicating your results to others, so we include our thoughts about how to write good papers and how to give good talks. A video presentation of material related to this chapter is [Freeman2020].

Many graduate students we know feel it is important to co-author many papers. While the number of papers a student has can be a rough measure of research productivity, it is our experience as mentors, faculty search committee members, and industrial research managers that *only the good papers count*. This is summarized in Fig. 42.1, a plot of made-up data that summarizes our impression, formed over many years, of the relationship between paper quality and its impact on one's career. Only the really good papers matter for your career. So it makes sense to learn how to write the best possible paper from any given piece of research work.

42.2 Organization

The PhD thesis advisor of two of us, Prof. Edward Adelson, wrote some good advice in response to a student's question of how to write a good paper [Adelson1992]:

1. Start by stating which problem you are addressing, keeping the audience in mind. They must care about it, which means that sometimes you must tell them why they should care about the problem.
2. Then state briefly what the other solutions are to the problem, and why they aren't satisfactory. If they were satisfactory, you wouldn't need to do the work.
3. Then explain your own solution, compare it with other solutions, and say why it's better.
4. At the end, talk about related work where similar techniques and experiments have been used, but applied to a different problem.

That structure fits well with a progression of section headings typical of many conference papers. As an example, here are the headings from paper [Fergus et al.2006], co-authored by one of us. The structure is useful for many research papers:

1. Introduction
2. Related work
3. Main idea

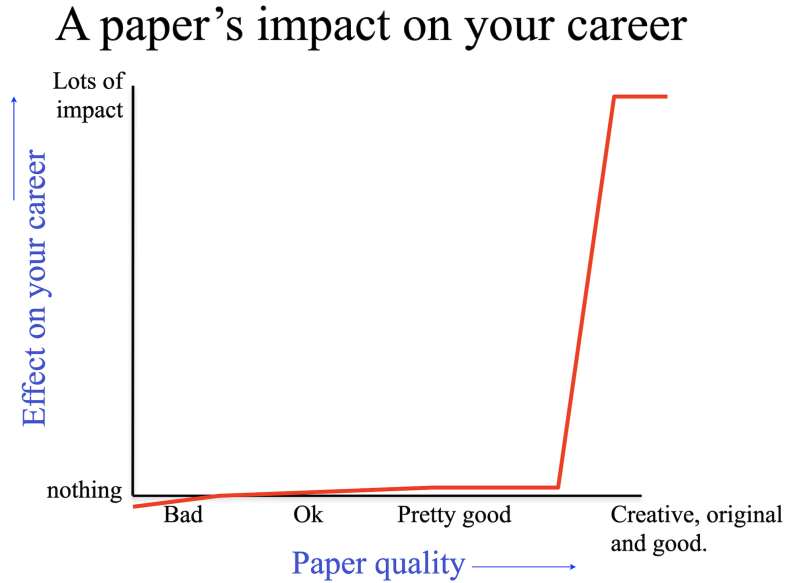


Figure 42.1: Plot showing the impact of a paper on one's career, as a function of paper quality. While the curve is subjectively drawn, it represents our judgment from many years of experience in research

4. Algorithm

- Estimating the blur kernel
 - Multi-scale approach
 - User supervision
- Image reconstruction

5. Experiments

- Small blurs
- Large blurs
- Images with significant saturation

6. Discussion

42.2.1 Introduction

Regarding paper introductions, Kajiya [Kajiya1993] writes, “You must make your paper easy to read. You’ve got to make it easy for anyone to tell what your paper is about, what problem it solves, why the problem is interesting, what is really new in your paper (and what isn’t), why it’s so neat. And you must do it up front. In other words, you must write a dynamite introduction.”

42.2.2 Main Idea

When appropriate to the paper, it can be very helpful to show a simple example that captures the main idea of the paper. Here is a figure from a different paper [Simoncelli et al.1992] that conveys the main idea of the paper very simply. The paper’s main idea was that wavelets lacked important desirable features for an image representation. Fig. 42.2 shows the failure of a wavelet representation to form a translation invariant representation. The top row shows two versions of the same signal, differing only in a 1-sample translation. The

bottom three rows show the coefficients of the high, mid, and low frequency bands of a wavelet representation of that signal. The figure points out in a simple way the drawback of a wavelet representation with aliased subbands that was the focus of the paper: as the signal translates, the signal representation energy moves to different frequency bands, and changes form within the mid frequency band.

42.2.3 Experimental Results

In earlier days of computer vision research, good ideas could be presented with only plausibility arguments to support them. However, current publication standards require that assertions in a paper be backed-up with experimental evidence. Readers will be expecting a table quantitatively comparing the performance of your algorithm with that of the previous state-of-the-art (SOTA). If you are proposing a new task for which there aren't previous benchmarks, you should still compare with an algorithm that solves a related task, noting that the algorithm was designed for a different task.

42.2.4 How to End the Paper

A cautionary note on how not to end the paper. Conference papers are often written under deadline pressure, and there is inevitably a list of results that the authors wanted to obtain that have been left uncompleted. You should resist the temptation to describe the items that were left unfinished in a "Future Work" section. It is hard to imagine a weaker way to end the paper than to enumerate all the things the paper doesn't accomplish—a summary of where it falls short.

Instead, the paper can finish with a review what has been presented, emphasizing the contributions. How has the world changed, now that the reader has read this paper? What new research directions have been opened-up; what can we do now that we couldn't do before?

42.3 General Writing Tips

Donald Knuth, the author of the classic book series, *The Art of Computer Programming*, wrote [Knuth et al.1989], "Perhaps the most important principle of good writing is to keep the reader uppermost in mind: What does the reader know so far? What does the reader expect next and why?"

My mental image is that of a great host, who anticipates every need of their house guest: after you arrive, they say, "Can I take your coat? Would you like something to drink?" At every moment, they know what you might be wanting and how to help you find it

42.3.1 Use Fewer Words

In **The Elements of Style**, Strunk and White [Strunk and White1999] write "Vigorous writing is concise. A sentence should contain no unnecessary words, a paragraph no unnecessary sentences, for the same reason that a drawing should have no unnecessary lines and a machine no unnecessary parts."

To make this point in the context of scientific writing, below is a paragraph that one of the authors of this textbook co-authored with a student.

The underlying assumption of this work is that the estimate of a given node will only depend on nodes within a patch: this is a locality assumption imposed at the patch-level. This assumption can be justified in case of skin images since a pixel in one corner of the image is likely to have small effect on a different pixel far away from itself. Therefore, we can crop the image into smaller windows, as shown in Figure 5, and compute the inverse J matrix of the cropped window. Since the cropped window is much smaller than the input image, the inversion of J matrix is computationally cheaper. Since we are inferring on blocks

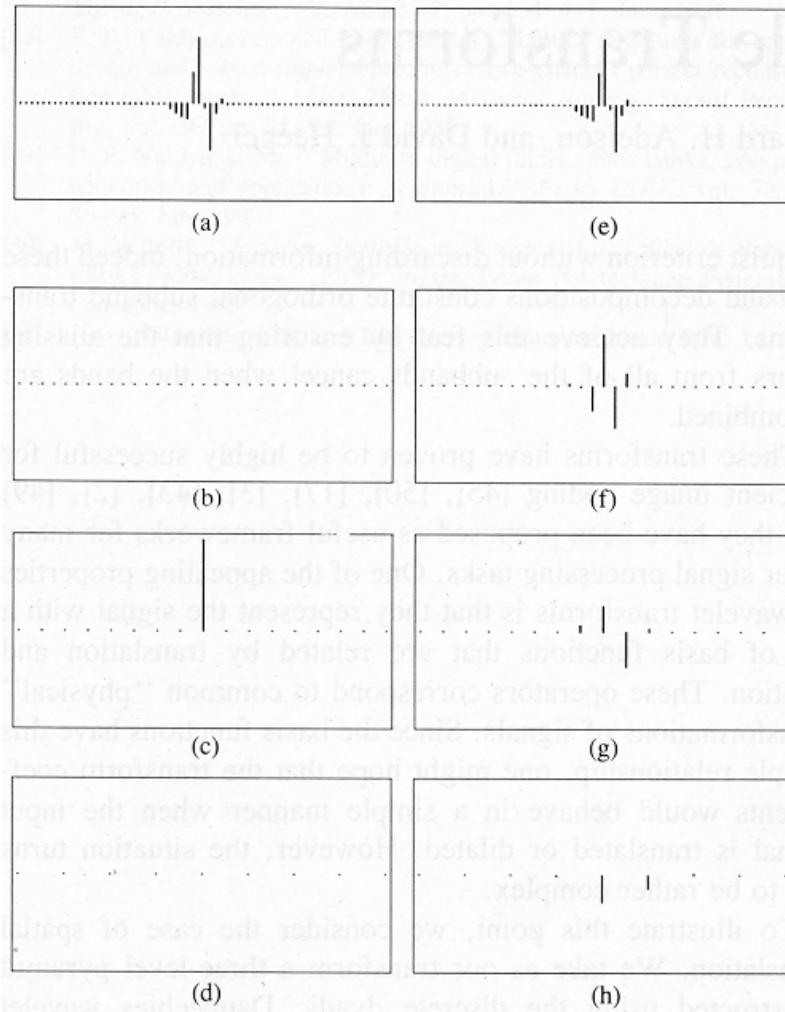


Figure 42.2: Effect of translation on the wavelet representation of a signal, caption and figure reprinted from [Simoncelli et al.1992]. (a) An input signal. (b), (c), and (d) show its high, mid, and low-frequency coefficients within a wavelet subband decomposition. (e) shows the same input signal, translated one sample to the right. (f), (g), and (h) show the decomposition of the shifted signal into the same three bands. Note the drastic change in the transform coefficients induced by the shift.

Algorithm	Dataset 1	Dataset 2	Dataset 3	Dataset 4	Dataset 5	Dataset 6	Dataset 7
Algorithm A	0.85	0.72	0.91	0.68	0.89	0.75	0.82
Algorithm B	0.78	0.81	0.88	0.79	0.85	0.83	0.80
Algorithm C	0.82	0.75	0.87	0.71	0.86	0.77	0.84
Algorithm D	0.80	0.73	0.89	0.70	0.87	0.76	0.81
Algorithm E	0.83	0.74	0.90	0.69	0.88	0.74	0.83
Algorithm F	0.81	0.76	0.86	0.72	0.84	0.78	0.80
Algorithm G	0.84	0.71	0.88	0.67	0.86	0.73	0.81
Algorithm H	0.82	0.74	0.89	0.70	0.87	0.76	0.81
Algorithm I	0.80	0.73	0.88	0.69	0.86	0.74	0.82
Algorithm J	0.83	0.75	0.87	0.71	0.85	0.77	0.80
Algorithm K	0.81	0.72	0.86	0.68	0.84	0.75	0.81
Algorithm L	0.82	0.74	0.88	0.70	0.86	0.76	0.80
Algorithm M	0.84	0.73	0.89	0.69	0.87	0.74	0.82
Algorithm N	0.81	0.75	0.86	0.71	0.84	0.77	0.80
Algorithm O	0.83	0.72	0.88	0.68	0.86	0.74	0.81
Algorithm P	0.80	0.74	0.87	0.69	0.85	0.75	0.80
Algorithm Q	0.82	0.73	0.88	0.70	0.86	0.76	0.81
Algorithm R	0.84	0.71	0.89	0.67	0.87	0.74	0.82
Algorithm S	0.81	0.75	0.86	0.71	0.84	0.77	0.80
Algorithm T	0.83	0.72	0.88	0.68	0.86	0.74	0.81
Algorithm U	0.80	0.74	0.87	0.69	0.85	0.75	0.80
Algorithm V	0.82	0.73	0.88	0.70	0.86	0.76	0.81
Algorithm W	0.84	0.71	0.89	0.67	0.87	0.74	0.82
Algorithm X	0.81	0.75	0.86	0.71	0.84	0.77	0.80
Algorithm Y	0.83	0.72	0.88	0.68	0.86	0.74	0.81
Algorithm Z	0.80	0.74	0.87	0.69	0.85	0.75	0.80

Figure 42.3: Blurred for anonymity, here is a prototypical table of results. Rows are different algorithms, columns are datasets or tasks, and the numbers indicate performance, with the best performances given in bold.

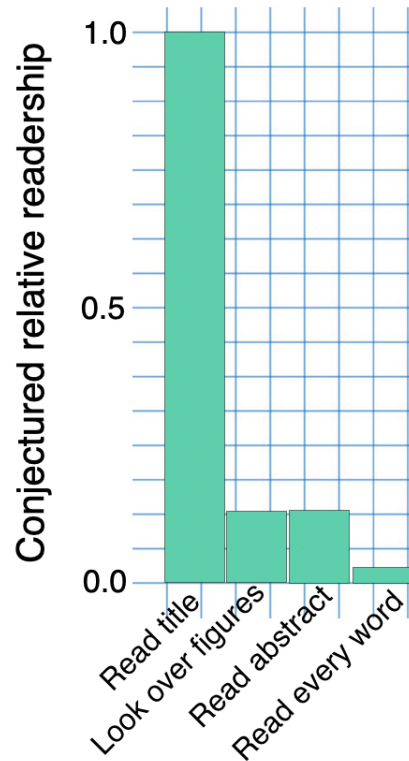


Figure 42.4: The readership of a computer science paper (conjectured relative amounts). Many people will read the title, but far fewer will read the abstract or look at the figures. Fewer still will read every word of the paper. You should let your readers learn something from your paper at each of those different levels of engagement.

of image patches (i.e. ignoring pixels outside of the cropped window), the interpolated image will have blocky artifacts. Therefore, only part of $xMAP$ is used to interpolate the image, as shown in Figure 5.

While the text above conveys the desired technical information, it describes things in a round-about way. The word count is 149. Below, the paragraph has been re-written to just 88 words without changing its meaning:

We assume local influence—that nodes only depend on other nodes within a patch. This condition often holds for skin images, which have few long edges or structures. We crop the image into small windows, as shown in Fig. 5, and compute the inverse J matrix of each small window. This is much faster than computing the inverse J matrix for the input image. To avoid artifacts from the block processing, only the center region of $xMAP$ is used in the final image, as shown in Fig. 5.

Note that writing more concise text helps your paper in two ways. First, one is typically fighting against a page limit, particularly in conference submissions, so tighter text lets the author say more, or add another figure. Second, concise text is usually clearer and easier to understand.

42.3.2 Title, Figure Captions, and Equations

While we may imagine our reader reading every word of our paper, in our current era of many, many published papers, most of your readers will probably only skim your paper. Based on subjective observations, we suspect that the engagement of readers with your paper will look something as shown in Fig. 42.4.

If only because that is where most of your readership lies, you should write your paper




Figure 3: (a) Time-frame assignments for the front-most surface pixels, based on stereo depth measurements alone, without MRF processing. Grey level indicates the time-frame assignment at each pixel. (b) Shape-time image based on those assignments. (c) Most probable time-frame assignments, computed by MRF. (d) Resulting shape-time image. Note that the belief propagation in the MRF has removed spurious frame assignment changes.

Figure 42.5: Figure and caption from [Freeman and Zhang2003]. The caption should be understandable by itself, and should point to what the reader should take away from the figure.

so that readers will learn from it, even at those reduced levels of engagement. The title should convey the top-level message of the paper. The figures and their captions should be self-contained and tell the story of the paper. The abstract should provide a good summary.

A good title can generate attention and interest for the paper. “What makes Paris look like Paris” [Doersch et al.2012] is a delightful paper, and the title lets the reader see that right away. In contrast, one of us co-authored a paper called “Shiftable Multi-scale Transforms”. While that title is appropriate, it’s not catchy. After the paper was in print, we realized that we should have added the subtitle, “What’s wrong with wavelets?”, a catchy phrase that tells the main point of the paper.

Because you can assume that many of the readers of your figure captions will not be reading the full paper, *the captions should be self-contained*. They should point out to the reader what they are supposed to notice about the figure, to help both the full-paper reader and the figures-only reader. Figure 42.5 shows an example.

Equations interposed with text can pose a challenge to the reader. Does the reader stop to work through each equation, or continue reading as quickly as if the equation were words? Knuth [?] and Mermin [Mermin1989] both give advice on writing with equations.

Knuth writes, “Many readers will skim over formulas on their first reading of your exposition. Therefore, your sentences should flow smoothly when all but the simplest formulas are replaced by ‘blah’ or some other grunting noise.”

Mermin adds, “*When referring to an equation identify it by a phrase as well as a number.* No compassionate and helpful person would herald the arrival of Eq. (7.38) by saying ‘inserting (2.47 and (3.51 into (5.13)...’ when it is possible to say ‘inserting the form (2.47 of the electric field E and the Lindhard form (3.51) of the dielectric function ϵ into the constitutive equation (5.13)...’ ”

42.3.3 Tone

The tone of your writing is important. You may feel pressure to over-sell, hide drawbacks, and disparage others’ work. It is in both your short-term and your long-term interest not to succumb to that pressure. If your work has a shortcoming, it is much better that you point it out than to have someone else do so later.

Papers written from the point of view that “we’re all good researchers, doing our best” are much more pleasant to read than papers written from the thesis that everyone else is an idiot. Note the delightful sentences written by Efros [Efros and Freeman2001] when he was a graduate student, “A number of papers to be published this year, all developed independently, are closely related to our work. The idea of texture transfer based on variations of [6] has been proposed by several authors (in particular, see the elegant paper by Hertzmann et.al. in these proceedings). Liang et al propose a real-time patch-based texture synthesis method very similar to ours. The reader is urged to review these works for a more complete picture of the field.” This generous tone wins over the reader, and promotes long-lasting friendships with people who study the same problems you do.

You should develop a reputation for being clear and reliable. Always convey an accurate impression of the performance of an algorithm. If something doesn’t work in important cases, let the readers know. In [Fergus et al.2006], we noted that our deblurring algorithm didn’t work for a case where we had been sure it should work, for MAP estimation of the deblurred image. Some years later, we learned, and described in [Levin et al.2011], that we had been wrong to believe that MAP estimation of the deblurred image was feasible. We were glad we had described our counter-intuitive result years before.

42.3.4 Author List

The quality of the papers you write matters so much more than your positions within the author lists of the papers. If adding more authors will make a paper better, then you should add more authors. If a collaborator feels they should be an author, but you yourself are not sure, we find it is generally a good idea to trust the collaborator and to add them as a co-author. It’s much better to be one of many authors on a great paper than to be one of just a few authors on a mediocre paper.

42.3.5 Avoiding Rejection

With the task of rejecting at least 3/4 of the submissions, conference area chairs, who make the acceptance decisions, are groping for reasons to reject a paper. Here’s a summary of reasons that are commonly used, and it is best not to provide the area chair with any of these easy reasons to reject your paper:

- Do the authors not deliver what they promise?
- Are important references missing (and therefore one suspects the authors not up on the state-of-the-art for this problem)?
- Are the results too incremental (too similar to previous work)?
- Are the results believable, with sufficient tests to support the claims?



(a) the cockroach



(b) Puppy with six toes

Figure 42.6: Two categories of borderline papers from an area chair’s point of view. (a) The cockroach is an unexciting paper which is difficult to kill. There’s nothing obviously wrong with it, but it is not exciting [Fashion2021]. (b) The puppy with six toes is a delightful paper, but with some clear flaw. The flaw may not be important, like a puppy having six toes, but it’s an easy reason to reject the paper [Hospital2021].

- Is the paper poorly written?
- Are there mistakes or incorrect statements?

As an area chair, the very good and the very bad papers are easy to deal with, and most of the decision-making effort goes to the borderline papers. One of us writes, “I find that much of my time is spent dealing with two kinds of borderline papers: cockroaches and puppies with six toes.

“Cockroaches” have no exciting contributions, but the reviews are ok and the results show incremental improvement to the state of the art. They’re hard to kill (hence the nickname), and maybe $2/3$ are accepted as posters, and $1/3$ are rejected. As an author, it’s best to work harder for the fresh results, to bring papers like this out of the cockroach category.

At the other end of the spectrum of borderline papers are “puppies with six toes”. These are delightful papers which easy-to-see flaw. The flaw may not matter (like six toes on a puppy), but makes the paper easy to reject, even though the paper may be fresh and wonderful. Maybe $2/3$ of those papers get rejected, and $1/3$ maybe accepted as posters. As an author, it may be better to wait another submission cycle with a paper like this, to remove the easy to spot flaw. Then the paper might receive an oral presentation at a subsequent conference submission cycle.

Chapter 43

How to Give Talks

43.1 Motivation

Giving good talks is important for a researcher. One might think, “The work itself is what really counts. Giving the talk is secondary”. But the ability to give a good talk is like having a big serve in tennis—by itself, it doesn’t win the game for you, but it sure helps. And the very best tennis players all have great serves. So you should learn to give good talks about your research.

There are many sources to go to for giving talks, and you should read or view several to pick out the tips that work best for you. Prof. Patrick Winston’s talk about giving talks is here [Winston2016]. Here are a few other good sources [Hoogterp2014]*****. Following are a set of heuristics that have worked well for us.

43.2 Prepare

Our top level message for how to give a good talk is the same as our message for how to not be nervous when giving a talk: “Prepare and practise”.

Regarding preparation: think through your talk. It’s almost always better to say a talk from notes than to read it from a script—the audience feels you’re speaking to them so it’s easier for them to pay attention. But for people giving a talk in a language that they are not comfortable with, writing out a script ahead of time may be a better solution. There may be parts of the talk that you feel are difficult to get through and for those parts, even for a native speaker, it may help to write out verbatim what you plan to say. When you give the talk, you might ignore the script, but having written it out once will make it easier to say.

When you can, you should visit the room where you’ll be giving the talk to identify any issues that may come up. Will you be blocking anyone as you stand at the lectern to give the talk? Will there be anyone to help you set-up for the talk? Before the talk, you should decide how to position yourself so that you can see your presentation, but also engage with the audience. Other things to check before the talk: Check your tools: is the microphone where you want? Are you occluding the screen while talking? Practice your voice volume.

The second great cure for being nervous about a talk is to practise it. Practise by yourself. Give practise talks to your friends. Think through the talk and find the story for how one part relates to the next. If you can’t find that story, you may want to reorder the talk until it flows naturally from one topic to another.

43.3 Nervousness

I used to be a nervous public speaker. I still find the hours before giving a lecture at a conference or to a class, to be stressful and I’m always reviewing my talk to try to improve

blah blah blah blah blah blah blah So that gives us the objective function we want to optimize. Now, how do we find the optimal value? There are two approaches you can take. blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah So now, with these tools in hand, we can apply this methods to real images. blah blah blah blah blah blah blah blah blah blah” You want to help even the distracted listener follow something of what you are saying.

You can also make your talk easier to follow if you add dynamics—variations in speed or intensity. I like to find a part of the talk that I feel is particularly magical or wonderful, and let that show through. When I gave talks about image deblurring, I would emphasize one aspect of it that I particularly love: “I love this problem; it’s beautifully underdetermined. There are many ways we can explain a blurry image. It could be that that’s what was there in the world— we took a sharp picture of a world that happened to look blurry. Or we took a blurred image of a sharp world.”

The audience loves to watch you be excited.

People Love to See a Good Fight

As described by Adelson [Adelson1995], the audience loves to watch a good fight. As a technique to help the audience follow what you’re presenting, you can set-up a fight between two competing conjectures. For example, you might say, “The flat earth theory predicts that ships will appear on the horizon as small versions of the complete ship. Under that theory, you’d expect approaching ships to appear as in Fig 43.1 (a). Whereas the round earth theory predicts that the top of the sails will appear first, then gradually the rest of the ship below it, resulting in approaching ships looking as in Fig 43.1 (b). Then they are waiting with anticipation when you show them the evidence from your experiment, Fig 43.1 (c).




Figure 43.1: (a) The appearance of ships approaching a port, from the flat-earth theory. (b) Their appearance as predicted by the round-earth theory. (c) The experimental evidence: a photograph of ships approaching a port □.

Show Yourself to the Audience

The actor Alan Alda describes the importance of engaging and connecting with the audience [Alda2014]. He noted the improvement in scientific speaking of volunteers after they engaged in improvisational theater exercises. The volunteers were primed to engage and connect with others, and it showed through in their scientific talks.

What the Audience Wants From a Talk

Here’s what I think an audience wants from a technical talk:

- To have everything follow and make sense
- To learn something
- To connect with the speaker, to share their excitement.

- They want to watch you love something!

If you have these elements in your talk, people may come up to you and say how much they enjoyed your talk.

43.5.3 How to End the Talk

How should you end your talk? Let's start with how not to end the talk: by asking, "Are there any questions?" The audience wants to clap at the end of a talk, but if the speaker closes by asking for questions, they don't know whether to stop, or to ask a question, and you may receive some light, embarrassed applause before someone asks a question [Adelson1995].

Better is to close with "Thank you." The audience knows that is their signal to applaud, and they will. Then you can ask for their questions.

We note that Patrick Winston disagrees with this [Winston2016]. He didn't like the implication that the audience was doing you some favor for attending your talk, and preferred to remind the audience of how he has kept his earlier promise to them regarding what they would learn after they listened to his talk.

Bibliography

- [Adelson1992] Adelson, E. H. 1992. personal communication.
- [Adelson1995] Adelson, E. H. 1995. personal communication.
- [Alda2014] Alda, Alan. 2014. Alan Alda on Improvisation for Communication of Science. <https://www.youtube.com/watch?v=j4XgjkXDxss>. MIT McGovern Institute talk.
- [Doersch et al.2012] Doersch, Carl, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A. Efros. 2012. What makes paris look like paris? *ACM Transactions on Graphics (SIGGRAPH)* 31 (4): 101–11019.
- [Efros and Freeman2001] Efros, A. A., and W. T. Freeman. 2001. Image quilting for texture synthesis and transfer. In *Acm siggraph*. In *Computer Graphics Proceedings, Annual Conference Series*.
- [Fashion2021] Fashion, Amazon. 2021. Fun World Adult Cockroach Costume. <https://www.amazon.com/Fun-World-Costumes-Cockroach-Costume/dp/B0038ZQYRC>.
- [Fergus et al.2006] Fergus, R., B. Singh, A. Hertzmann, S. Roweis, and W. T. Freeman. 2006. Removing camera shake from a single image. In *Proc. siggraph*.
- [Freeman and Zhang2003] Freeman, W. T., and H. Zhang. 2003. Shapetime photography. In *Proc. ieee computer vision and pattern recognition*.
- [Freeman2020] Freeman, William T. 2020. How to write good papers. starting 9:00 in: <https://www.youtube.com/watch?v=W1zPtTt43LI>. in CVPR 2020 workshop, How to Write a Good Review, organized by Laura Leal-Taixé and Torsten Sattler.
- [Gilbert] Gilbert, Rob. Quote. <https://www.quotes.net/quote/13310>.
- [Hoogterp2014] Hoogterp, W. 2014. *Your perfect presentation*. McGraw-Hill Education eBooks.
- [Hospital2021] Hospital, San Jose Animal. 2021. <http://www.sanjoseanimalhospital.com/puppy-and-kitten-packages>.
- [Kajiya1993] Kajiya, Jim. 1993. How to Get Your SIGGRAPH Paper Rejected. <https://www.siggraph.org/sites/default/files/kajiya.pdf>. SIGGRAPH 1993 Papers Chair.
- [Knuth et al.1989] Knuth, D. E., T. L. Larrabee, and P. M. Roberts. 1989. *Mathematical writing*. Mathematical Association of America Notes. Series Number 14.
- [Levin et al.2011] Levin, A., Y. Weiss, F. Durand, and W. T. Freeman. 2011. Efficient marginal likelihood optimization in blind deconvolution. In *Proc. ieee computer vision and pattern recognition*.

- [Mermin1989] Mermin, N. David. 1989. What's wrong with these equations? *Physics Today*.
<http://www.ai.mit.edu/courses/6.899/papers/mermin.pdf>.
- [Simoncelli et al.1992] Simoncelli, E. P., W. T. Freeman, E. H. Adelson, and D. J. Heeger.
1992. Shiftable multi-scale transforms. *IEEE Trans. Info. Theory* 2 (38): 587–607.
- [Strunk and White1999] Strunk, William, and E. B. White. 1999. *The elements of style*.
Allyn and Bacon : Longman.
- [Winston2016] Winston, Patrick. 2016. How To Speak. <https://vimeo.com/101543862>.

□ .