Chapter 23

Lecture 23: Project papers and talks

Monday, May 2, 2011 MIT EECS course 6.869, Bill Freeman

Your final projects involve, in addition to doing the actual reading and experimentation, writing a 5-8 page paper describing your work, and presenting a short, 5 minute, talk. To help you with that, today's lecture is about writing papers and giving talks. Since almost all of you are involved with academic research, or will be, we decided to broaden out this lecture and have it also pertain to writing papers and giving talks in general academic research, not just for this class. What makes a conference paper good will also your projects be good, and we hope the general academic advice will be useful to you.

23.1 Papers

Let's start with writing papers. First of all, when should you write a paper? You're doing your research, making progress, when do you decide, "ok, I'd better package this up now"? I think in some ideal world, it would be entirely driven by the progress of the work itself. You get a good result, you write it up, and you submit it somewhere.

But, in practice, paper submission is often driven by conference or special issue deadlines. This can be both a good and a bad thing. A conference deadline can cause people to submit something before it's ready. On the other hand, these deadlines are great forcing functions, and lots of good work can happen in the month preceeding a conference submission deadline.

In deciding when to write up research work, an important thing to note is the following (subjective) curve, showing the benefit to one's career for publishing a paper as a function of the quality of that paper. The curve is flat (or below zero!) for almost all papers until the quality level becomes very good, and then a paper is very valuable. My point: only the good papers count. Nobody remembers boring papers, and nobody cares about them. It probably wasn't worth the effort to write them up.

Packaging up your results and writing up a paper is a lot of work, and there's a real opportunity cost to publishing mediocre papers. You may miss out on aiming high and getting some really good result.

My own experience: when I was a graduate student, I only wrote up things when I felt I had something to say. I remember two conference paper submission deadlines when the lab was in a frenzy of paper-writing and I was just doing my research because I didn't have a good result then. That turned out to be a good strategy for me.

Let's talk for a bit about where to publish papers. For the past 10 or 20 years, the "action" in computer vision is at the top conferences. That's where the exciting new results are published and read. Journals are still useful, but they play a different role now. Now they're used for coherent summaries of several

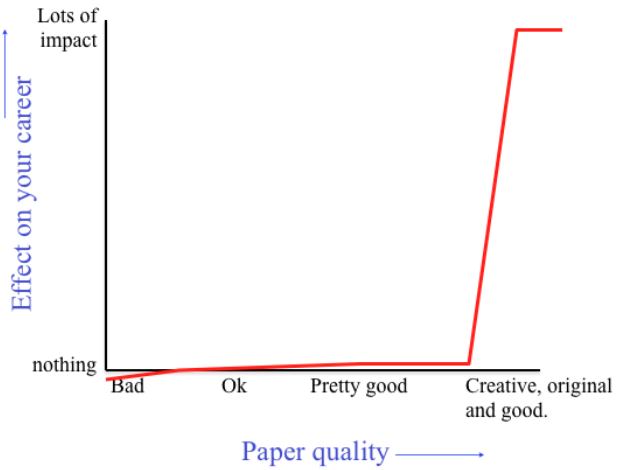

Figure 23.1: Curve showing a paper's impact on your career as a function of its quality.

conference papers, or for making archival versions of your favorite conference papers. Journal papers allow for more detail, and more perspective in an article.

There are 3 top-tier conferences in computer vision:

- CVPR, Computer Vision and Pattern Recognition, annual conference, always somewhere in North America, but very international in membership.
- ICCV, International Conference on Computer Vision, every two years, cycles through being in North America, Europe, and Asia.
- ECCV, European Conference on Computer Vision, happens every year that ICCV does not happen.

All these conferences have 20-25% acceptance rates, and awards for the best papers, etc.

Oftentimes computer vision research borders on computer graphics research, or machine learning, so important related conferences are:

- SIGGRAPH, stands for Special Interest Group on Graphics. This is the top computer graphics conference. Since computer graphics now uses captured images and surfaces, they often publish work using computer vision or image processing techniques.
- SIGGRAPH Asia, is a version of Siggraph in Asia. Almost as prestigious as Siggraph, but not quite.
- NIPS, Neural Information Processing Systems. The top machine learning conference, publishes work on applications of machine learning methods to vision, or relating to the statistics of natural images.

SIGGRAPH is such a high-quality publication venue that simply counting the number of Siggraph publications is a reasonably good measure of impact in graphics.

The journals, Science and Nature, are also very prestigious publication venues, although computer vision work is seldom published there. The work needs to have a science component.

23.1.1 a CVPR paper

Ok, so let's assume you're submitting it to CVPR. This is exactly the format of paper we'd like to see for your final projects, so now all the rest, regarding writing, applies equally well to your final projects as to any research project you may be submitting to CVPR. The paper can be eight or fewer pages. Here is a template for the typical organization of a paper:

- Introduction
- Related work
- Image Model [[main idea]]
- Algorithm
- Experiments
- Discussion

These headings (and the images) are from a paper of mine. I'm not saying it's a great paper, but it's a well-organized paper.

Introduction

Now let's examine the content, what to say within these headings. My thesis advisor, Ted Adelson, had good advice regarding writing a paper. He gave this in response to an informal survey by a graduate student.

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 bettter.
4. At the end, talk about related work where similar techniques and experiments have been used, but applied to a different problem.

Since I developed this formula, it seems that all the papers I've written have been accepted. (told informally, in conversation, 1990).

Most of those points would be made in the introduction and related work section of the paper. Typically, the Introduction tells the story using a broad brush. Sometimes the professor may write the introduction, since they know best how to situate the work relative to what other people have done, while the graduate student may write the rest of the paper, since they know the details of what they did. Some of my friends complain that all they write now are introductions. The reader should have the high level story in mind, after they read the introduction–the problem the paper addresses, the paper's main contribution or insight, and a rough idea of the techniques used.

Related work

The related work section is easy to write poorly. There is a temptation to recite a list of authors, with a one or two sentence description of each group's approach after their names. What the reader really wants is a description by contrast–what was the innovation of their work, and how does it still fall short. (if it didn't fall short in some way, there would be no reason to write the current paper). That type of description is much easier to read, as everything is written in relation to the current paper.

Main idea

After the related work, one often describes the core idea or insight of the paper. A technique that I don't think is used enough is to explain the main idea by way of a simple, or toy, example. Such examples in doing the research itself. This example is from our "Shiftable multiscale transforms" paper. We wanted to make the point that non-over-complete wavelet transforms don't give you a translation invariant representation of the signal energy in different frequency bands. The top row is two different signals, identical except for a 1 pixel offset between. In the three rows below is the signal representation in 3 bands of a wavelet transform. On the left, all the signal energy is represented as a delta function at one wavelet position and scale. But shift that signal over by one pixel, and–blat–the energy gets dispersed into the other frequency bands. (That's because each band is an aliased representation of the signal. These aliased subbands rely on information from other aliased bands to reconstruct the signal. The details of those aliasing cancellations depend on the signal position).

The point here is not about the wavelets, but about using a simple example to reveal an insight of the paper in an easy-to-understand way.

Experimental results

In the old days, people published many speculative computer vision papers. It was like the wild west. There was something exciting about that time, but it's gone. Now, any innovation must be accompanied by an end-to-end emperical study validating the proposed algorithm. We lose a little bit of creativity with that change, but we gain much-needed rigor.

Another innovation: it is now customary, and almost required, to release computer code or data with every publication. Needless to say, this is a huge benefit to the field.

How to end the paper

This paper I've been showing you doesn't end with a conclusions section; it ends with "discussion" (we ran out of space). Sometimes papers end this way, but usually you end with a Summary or Conclusions section. You can reiterate the main points, and restate the paper's conclusions. Ideally, you'd like to state the results in a new, stronger way. Now that the reader has read the paper, can you state the results more powerfully? How can this paper change how people will do things in the future?

How not to end the paper

Apparently, after writing a paper, there is often an urge to list everything that the authors had wanted to do but couldn't get to work. This is listed in a section called "Future work". I can't think of a weaker way to end a paper than to provide an itemized list of everything that was expected to be easy with this algorithm, but which didn't work out. There's no partial credit given for having an idea that would have been really cool if it had worked out.

It is ok, however, if this paper opens up new research directions, to close by describing those new directions. But that's different than describing things you didn't get around to doing.

23.1.2 How conference papers are evaluated, and what that means for you as an author

Now let's talk about how conferences are run, and how papers are evaluated. Two or three conference program co-chairs are appointed or elected by the conference, or by a committee representing the conference. (I was the program co-chair for ICCV in 2005, and, even though I vowed not to do it again, will be for CVPR in 2013). The program co-chairs need to have integrity, know the field, and know who the good people are in the field. The co-chairs select 40 or so people to be "area chairs"–they handle all the papers within a particular area, and make the final decisions on each paper.

After the 1600 papers are submitted, the program chairs, with the help of author keywords, decide which area chair will handle which paper. The area chairs then select 6 or 7 candidate reviewers for each of the 40 papers they will handle. An algorithm then addresses the NP-hard assignment problem of how to allocate reviewers to papers. You can imagine that the good reviewers have far more requests for their services than they can handle.

The reviewers read their 6-12 papers and rate and write reviews for each of them. The area chairs then try to encourage the reviewers of borderline papers and papers with divergent reviews to read the others' reviews and come to a concensus opinion about the paper. Some reviewers will modify their review in light of points made in the reviews of others. That removes some statistical independence of the reviews, but lets the reviews make more sense when read together by the authors.

The authors then read the reviews, and write a brief rebuttal to them. (In rebuttals: be positive, be respectful, be succinct.)

Finally, the area chairs meet at some central location and discuss the papers in small groups, and then in larger groups, to try to arrive at a uniform set of criteria and threshold for paper acceptance, and for the decision of poster or oral presentations.

23.1.3 General paper writing tips slide

Given this process, run by volunteer, rushed, falliable humans, how can you maximize the chances of your paper getting accepted? There are many things you can do, and, fortunately, they fit closely with what it takes to write a good paper.

Many of the quotes on the slides in this section are from an excellent paper by Jim Kajiya on how to write a good conference paper, called "How to get your SIGGRAPH paper rejected".

You want to convey to the reader the main point of your paper before you get bogged down in the technical details. So you have to write a great introduction. My image of how scholarship works had been something like medieval monks, pouring over each manuscript, with lots of time on their hands. That was my image of both reviewers, and the readers of published articles. But in reality, both reviewers and readers are folks with a lot of other things on their minds, like shoppers at a large bazaar. You have to make reading and understanding your paper as easy as possible for them.

Kajiya points out that you have to write a dynamite introduction. "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." This is a similar message to what Adelson said, but stated in a slightly different way. Knuth gives a similar emphasis on starting the paper well.

I like this quote from Knuth, "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?" I like to think of our task as writers as being good hosts to a guest in our house. A great host

will anticipate your needs, “here, let me take your coat. You’ve walked a long way–sit down. You must be hungry now; let’s eat”. What ordering of the material will make the most sense to the reader? What transitional sentences will let everything flow as logically as possible?

23.1.4 Tips and advice

Writing style

Equations Knuth and Mermin give a similar message regarding equations: number them all, but, more importantly, refer to them by a phrase, as well as by a number. Instead of writing, “Inserting (2.47) into (5.13)...”, you should write, “inserting the form (2.47) of the electric field E into the constitutive equation (5.13)...”. As Knuth writes, you should assume that some people will read your paper (at least, on the first read-through) with the word “blah” replacing every equation. You can’t assume they’ll make a second reading and you need to write the paper to account for that reader who says “blah” to every equation (they might be a reviewer!).

Editing At least half of writing is re-writing. I must confess that I hate wordy writing, and I think an important part of editing is scanning for sentences that can be made shorter and clearer.

Figures Ideally, the figures of the paper should convey the main points of the paper by themselves. A reader should be able to look through the figure, read the captions, and understand the main point of the paper. (Some of your readers will approach a paper that way, and you should write for them, as well as your other readers). Oftentimes, you’ll want to tell the reader in the caption what message you want them to take away from this figure. The caption of the example figure shown ends with “...Note that the belief propagation in the MRF has removed spurious frame assignment changes.”

Writing strategy

The points above are general advice for writing papers. The following points address the more specific task of writing a paper that gets accepted to the conference. Fortunately, they also improve the writing.

With 70-80% of submissions to the conference rejected, reviewers and area chairs are looking for reasons to reject a paper. You want to make sure none of the following questions can be answered the wrong way.

- Do the authors promise more than they deliver? More on this in the next slide.
- Are there some important references that they dont mention (and therefore theyre not up on the state-of-the-art for this problem)?
- Has their main idea been done before by someone else? If this criticism might apply to this paper, you need to make this paper’s contributions clear to the reader. It’s fine to build on others’ work, and you need to make it clear how you’re doing that.
- Are the results incremental (too similar to previous work)?
- Are the results believable (too different than previous work)? The tension between this point and the previous one can certainly be a source of frustration. If your results are phenomenal, you may need to work hard to persuade the reader that they are for real.
- Is the paper poorly written? Spend time re-organizing and re-writing the paper.

- Do they make incorrect statements? Having knowledgable friends read drafts of your paper can help catch statements that might cause problems.

Note that all the above concerns are surmountable, but they all take time to address. Be sure to leave yourself enough time to write a great paper. Ask other people to read your paper and give comments.

Other stylistic tips

Only promise what you deliver

A personal story relating to only promising what you deliver: About 10 years ago I wrote a paper with a student intern at MERL. As always, I had big plans for the research I hoped the student could accomplish by the ICCV submission deadline. I wanted to build a system that would learn local evidence for shading and reflectance, then propagate that local evidence in a Markov Random Field to find a global separation of shading from paint. But, again, as always, we fell short of that goal–despite a last-minute push, we didn't get the propogation working in time for the submission deadline.

So we wrote up what we had done, reducing the scope of the paper to only include what we had results for. We titled the paper, "Learning local evidence for shading and reflectance", and wrote the paper with that reduced goal in mind. It was accepted, and certainly wouldn't have, if we had described our initial goals as the goal of the paper. If that sounds disengenuous, note that change of scope also makes the paper easier to read and understand.

Be kind and gracious

In a draft of a paper I wrote as a graduate student, I said "Snickerdoodle et al didn't show their processing results for anything more complicated than a single line". My thesis advisor pointed out that some time in the future, Snickerdoodle will be sitting on some NSF panel, evaluating some proposal of mine, and you don't want him to have a grudge against me. I toned down my criticism of the work, which probably hadn't been an accurate characterization, anyway.

A champion of writing gracious text is Alyosha Efros, at CMU. He was a summer intern of mine at MERL, and he was very generous in describing concurrent work that was closely related to what we were publishing,

"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 [9, 1, 11] (in particular, see the elegant paper by Hertzmann et.al. [11] in these proceedings). Liang et.al. [13] propose a realtime 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."

Being generous conveys authority and comfort; it's the tone you want to have in your papers. I should point out that there was no loss in our being generous there–that paper from 10 years ago in which we praised our competitors is one of the most highly cited graphics papers of that decade.

Be scrupulously honest

There are perceived pressures to oversell one's work. Don't succumb! It is in both your short-term and your long-term interests to be clear and honest about how your algorithm works and about its limitations. In describing the many reasons for selecting a particular paper for the conference's best paper prize, I heard a conference program chair say, "and because the author was Schnabovsky, I knew I could trust it." That's the reputation you want to develop. Considering any time scale longer than about a week, both you lose and the field loses if you publish a paper implying that some technique works well when really it doesn't. It takes other researchers some time to figure out that the technique really isn't any good, and then you've lost your credibility, which can take a very long time to repair. (Of all the papers to read, why would people spend their time reading your papers if they don't trust you? And how can you regain peoples' trust if no one reads your papers?)

Authorship and author order (slide 50, 51)

You have some research that is pretty good. If you add the expertise of your officemate Finch-Liu, it could become a spectacular paper, but then you might not be first author. What should you do? Remember the quality-impact curve of Figure 23.1. It is much, much better to be second-author on a great paper than to be first-author on a mediocre paper.

Author order is not nearly as important as the research that gets done. When people evaluate you for a job, yes, they look at the papers in your resume, but they primarily look at your reference letters. In those letters, there is plenty of opportunity for people to spell out who did what on a paper, if that's important to clarify.

Paper title (slides 52, 53)

Paper titles do matter, and help determine which papers get read. In practice, it helps the fate of a paper if it has a memorable or catchy title. Same for an algorithm. For example, a recent algorithm for approximate nearest neighbor finding is called "Patch match"–wonderfully catchy. It has caught on.

An example from my own work–I co-wrote a paper that we called "Shiftable multi-scale transforms". Kind of a boring title, and not to informative by itself. As time went on, we realized we should have called it "What's wrong with wavelets", because that was the paper's main point. I think the paper would have had more success if we named it that.

the look of a CVPR paper

http://vision.ucsd.edu/sites/default/files/gestalt.pdf

By the way, there is a joke paper, called Paper Gestalt, that is worth looking at. They trained a classifier to look at visual features of a paper and to predict whether it would be accepted or rejected at the main vision conference, CVPR. Positive examples were from the CVPR proceedings. Negative examples were from the associated *workshop* proceedings of CVPR. They made a classifier correctly rejected 50 % of the workshop papers at the price of rejecting 15% of the accepted CVPR papers.

23.2 You, the corporation

As a researcher, you are a little corporation, and, like a corporation, you have different roles that must be played at different times. There is your research department, the part of you that actually does research. But you also need marketing, planning, and, as your career progresses, fundraising. Let's talk about the marketing role.

Doing good research is an essential task of a researcher, but describing the work to others is also essential. The ability to give good talks is important for a research career, much as having a big serve is important for a tennis professional. It's possible to play tennis without a good serve, and even possible to win sometimes, but all the top players have big serves, and, with very few exceptions, the top researchers give good talks.

23.2.1 Sources

There will be lots of ways to learn how to give good talks. Your advisor may coach you, friends may give you feedback if you ask for it. There are occasional talks on how to give good talks–Patrick Winston, a prize-winning teacher and former lab director, gives a great talk about giving talks every IAP, called Winston Speaks. Everyone will have their own recipe, and you should listen to them all and take from

them what makes sense to you. Based on listening to them, and my experience, here's what I think is important.

23.2.2 The most important thing

First, the high order bit–how to give good talks and not be nervous when giving a talk: prepare. Practice giving the talk, by yourself or to friends. If you want to be sure to get some part right, like the introduction, you may want to write out what you'd like to say, just to think through the progression you want to follow. Ahead of time, be sure to see where you'll be speaking. If you're speaking to a big group, I think it always helps to stand at the podium (during a break in the conference) and imagine speaking to the crowd. This always works for me: the best cure for nervousness is preparation.

23.2.3 3 different classes of talks

I find I give 3 different categories of talks:

- Very short talks, 2 to 5 minutes long
- Conference talks, 10 - 20 minutes long
- Seminar or invited talks, 40 - 60 minutes long.

Very short talks

You really have to rehearse these talks. In a long talk, you an afford to search for words, but you don't have the time in a very short talk. These may be in a fast-forward session in front of a very large audience, and you owe it to that large audience to give a clear, practised talk.

You want to state the problem, why it's interesting, and why it's hard. You'll want to say the key to your approach, and then show your results.

The goal is to convey the main idea, without getting lost in the details. At Siggraph, they have a "fast forward" session, when the entire conference is given, in 2 minute chunks, in about an hour. It is by far the most popular session of the conference, which I think tells us that even a 2-minute summary of a talk can be very valuable.

Longer talks

Conference length (10-20 minute) and seminar length (40-60 minute) talks share many characteristics, and I'll just talk about both of them together. Let's start with the bad news: the more you prepare for your talk, the better it will become. Why is that bad news? Because there is no end to it, the more you work on your talk, the better it will be.

What are some of the things to pay attention to as you work on your talk? You want to make it as easy as possible for the audience to follow your talk. Often, our image of our audience is that of a magical being. The audience is in rapt attention, listening to every word. In reality, the audience is often just as we ourselves are–perhaps tired or hungry, and not wanting to have to sit through another talk. You have to make your talk easy to understand even for someone like that.

How do you do that? Figure out how every part follows from every other. Understand how every part of the talk motivates the next. If a part doesn't seem to motivate the next, consider re-ordering the talk until it does.

Just as some readers will substitute the word "blah" whenever you write an equation, and you need to prepare for that, so will some listeners substitute "blah" whenever you recite any scientific details, and you need to prepare for those people, too. You can do that by layering your talk. Give periodic high-level summaries that tell the listener where we are in the story. These serve the function of captions in writing.

You become like a sea diver. You start at the surface, speaking in general terms, then you dive into the details. To give the listener a break and a chance to orient, you then summarize what that section was about, and transition to the next topic and discuss that.

ways to engage the audience

So youve been talking on and on. You want to break things up and keep the audience engaged. Can you think of a way to bring the audience into the talk? (ask the audience a question!)

Demos can also help. Since we work with vision, we can often show perceptual demos, which can do very well to bring the audience back into the talk. Present to the audience what the task is that the human or computer has to solve. That will at least get them thinking about things.

tell your story,

I think we have been designed to listen to other people tell the story of how they trapped the bear that had been chasing them and killed it. Or how they escaped from getting caught by the neighboring tribe. Find the drama, the anecdote, the personal story in what you have to tell people and bring that out. Maybe insert the people who did the science into the story about the science. It will make it so much easier for them to listen to your science story.

people like a good fight between competing explanations

People like to watch a good fight, and often a natural way to tell your research results is to frame it as a fight between two competing hypotheses. For example: "The *flat-earth theory* predicts that ships approaching port will appear as small versions of the ship, then gradually get larger and larger. The *round-earth theory* predicts that the top of the ship will appear first, then gradually the rest of it will appear below it. We did an experiment to distinguish between those two hypotheses, setting out a number of ships, and examining their appearance at varying distances from the port. Our results, showing the sails appearing first, support the round earth theory."

add dynamics

I don't like for talks to be at the same level of intensity throughout. As when you tell a story, you'll want to let yourself get excited in some parts.

I like to identify some part of the talk that really grabs me, and let that come through. For example, when I give talks on removing camera shake from photographs, I like to share with the audience what really excites me about that problem: the fact that it is wonderfully underdetermined. To explain a blurry image, it could be that the blurry image was really out there–that's what the world happened to look like right then–and you took a sharp photograph of it. (that's the explanation of the top row). Or it can be some crazy collection of intensities, blurred by a large Gaussian blur (middle row). Or it could be the sharp image that we're looking for, blurred by the kernel that we seek to get rid of (bottom row). This ambiguity forces you to come up with a statistical model for what images look like, and that's what I find so interesting about the deblurring problem–you have to think through what it takes to make one set of numbers look like an image, and another set not look like an image.

What the audience wants

I think the audience really wants to connect with you, and to hear your story. I found it helpful to listen to 10 minutes of the talk by Alan Alda about how he helped scientists give clearer talks by taking them though theater improvisation exercises. The exercises were designed to let you connect and empathize with your partner, and they help speakers connect with their audience. Then they talk with them, rather than lecture to them. Let yourself shine through.

How to end your talk

Don't say "are there any questions?", because then the audience doesn't know whether to clap or start asking you questions. Best is to say, "thank you", then they clap, then you or your host can ask for any questions, if there is time.