

Schedule

- Thursday, May 5:
 - Tracking humans, and how to write conference papers & give talks, **Exam 2 due**
- Tuesday, May 10:
 - Motion microscopy, separating shading and paint (“fun things my group is doing”)
- Thursday, May 12:
 - 5-10 min. student project presentations, **projects due.**

How to write a conference paper

Bill Freeman
MIT CSAIL
May 5, 2005

Sources on writing technical papers

- How to Get Your SIGGRAPH Paper Rejected, Jim Kajiya, SIGGRAPH 1993 Papers Chair, <http://www.siggraph.org/publications/instructions/rejected.html>
- Ted Adelson's Informal guidelines for writing a paper, 1991. <http://www.ai.mit.edu/courses/6.899/papers/ted.htm>
- Notes on technical writing, Don Knuth, 1989. <http://www.ai.mit.edu/courses/6.899/papers/knuthAll.pdf>
- What's wrong with these equations, David Mermin, Physics Today, Oct., 1989. <http://www.ai.mit.edu/courses/6.899/papers/mermin.pdf>
- Ten Simple Rules for Mathematical Writing, Dimitri P. Bertsekas http://www.mit.edu:8001/people/dimitrib/Ten_Rules.html

Why publish?

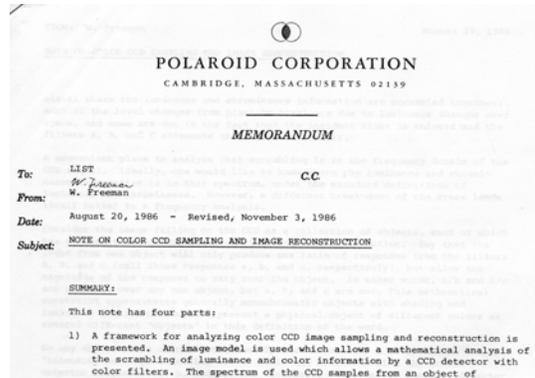
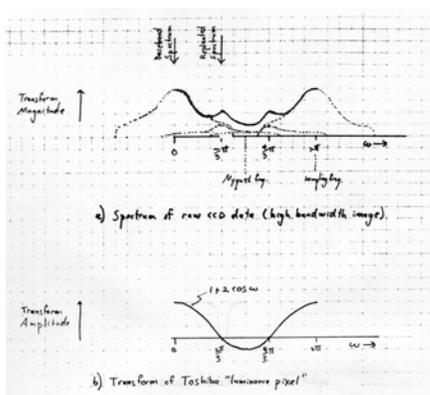


Figure from that memo...



Polaroid collaborated with Philips: a parallel universe!

PHILIPS GLOBAL HOME Let's make things better

YOUR COUNTRY CONSUMER PRODUCTS PROFESSIONAL PRODUCTS SEARCH

PHILIPS.COM About Philips News Center Careers Events Investor Info Theme Cache

Research

Philips' vision on research is to create technologies that will lead to products for improving people's lives. Its research capabilities are brought together in Philips Research, which was founded in 1914 in Eindhoven, the Netherlands.

Today, Philips Research is one of the world's major private research organizations with laboratories in the Netherlands, Belgium, England, France, Germany, Italy, the United States, China, Taiwan and India, collectively some 2,500 people.

A primary reason to publish:

To participate in the academic community

Some other reasons to publish

- To become well-known (to a very small group of people)
- To get more grant money
- To help get a job after graduation
- To publicize some product

Where publish

- Journal
 - Long turn-around time
 - But “archival”
 - Counts more in tenure decisions
 - Have a dialog with reviewers and editor.
- Conference
 - Immediate feedback
 - Publication within 6 or 7 months.
 - One-shot reviewing. Sloppier reviewing.

Kajiya on journal vs conference

“The emphasis on both speed and quality makes the reviewing process for SIGGRAPH very different from of a journal or another conference.

The speed and quality emphasis also puts severe strains on the reviewing process.

In a journal, the reviewer and authors can have a dialog where shortcomings and misunderstandings can be resolved over a leisurely pace. Also, even if there are significant flaws in a paper for another conference, the chances are that strengths will overcome the weaknesses in the judging.

In SIGGRAPH, if the reviewers misunderstand your paper, or if some flaw in your paper is found, you're dead.”

Special journal issues have some of the advantages of both



IEEE Transactions on Pattern Analysis and Machine Intelligence

Call for Papers

Special Issue on

Graphical Models in Computer Vision

IMPORTANT DATES

Paper Submission Deadline: 24 June 2002
Acceptance Notification: 24 January 2003
Final Manuscript: 24 February 2003
Publication Date: July 2003

By the way, I'm co-editing a special issue of IJCV on vision and learning, submission deadline of August 15, 2005.

CALL FOR PAPERS

Special Issue: Learning for vision and vision for learning.

Computational Vision and Machine Learning have become synergistic fields of research. Modern machine learning techniques have permitted large experimental improvements as well as a re-thinking of key problems such as recognition. On the other hand, vision has broadened the scope of machine learning offering rich and challenging new problems.

We solicit papers describing machine learning methods developed for or adapted to vision tasks and representations (and vice versa), such as

- priors and kernels useful for particular tasks
- machine learning algorithms addressing vision problems, e.g. fast detection, multi class categorization, semi supervised learning etc
- representations learned from images or videos, or optimized for visual inference

We wish to make the ideas and experiments presented in this special issue very easily accessible to other researchers.

We will therefore require all authors to:

a) Post their data (training and testing) on the web.

b) Make their code available in a form that allows other researchers

Some relevant conferences

- SIGGRAPH (ACM Special Interest Group on Graphics)
 - 350 submissions, 20% acceptance
 - Good, careful reviewing.
 - Some vision-and-graphics and learning-and-graphics.
- NIPS (Neural Information Processing Systems)
 - 300 submissions (?), ~25% acceptance
 - Reasonable reviewing.
 - Vision is a sidelight to the main machine learning show.
- CVPR/ICCV (Computer Vision and Pattern Recognition/Intl. Conf. on Computer Vision)
 - 700-900 submissions, 25-35% acceptance
 - Uneven reviewing
 - The main venues for computer vision and machine learning applied to computer vision.

Kajiya on conference reviewing

“The reviewing process for SIGGRAPH is far from perfect, although most everyone is giving it their best effort.

The very nature of the process is such that many reviewers will not be able to spend nearly enough time weighing the nuances of your paper. This is something for which you must compensate in order to be successful.”

Our image of the research community

- Scholars, plenty of time on their hands, pouring over your manuscript.



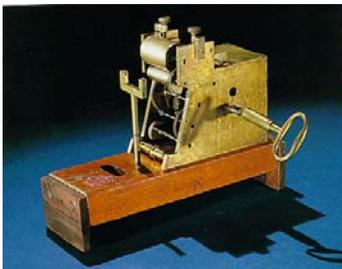
The reality: more like some large, outdoor bazaar



The conference paper review process

- Papers arrive (most on day of deadline)
- Conference chairs distribute papers to program chairs (20 – 60 papers to each person)
- Program chairs assign the papers to reviewers.
- 3 (NIPS, CVPR) to 5 (SIGGRAPH) reviewers read your paper.
- Program committee members meet to decide which papers to accept. The reviewers’ scores give an initial ranking; the program committee members then push papers up or down. NIPS: not much discussion. SIGGRAPH: lots of discussion.

How do you evaluate this complex thing, this paper?



Kajiya description of what reviewers look for

The most dangerous mistake you can make when writing your paper is assuming that the reviewer will understand the point of your paper. The complaint is often heard that the reviewer did not understand what an author was trying to say

Make it easy to see the main point

Your paper will get rejected unless you make it very clear, up front, what you think your paper has contributed. If you don't explicitly state the problem you're solving, the context of your problem and solution, and how your paper differs (and improves upon) previous work, you're trusting that the reviewers will figure it out.

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.

Kajiya description of what reviewers look for

Again, stating the problem and its context is important. But what you want to do here is to state the "implications" of your solution. Sure it's obvious....to you. But you run the risk of misunderstanding and rejection if you don't spell it out explicitly in your introduction.

Kajiya: well organized more important than well written

Really, you don't have to have a literary masterpiece with sparkling prose.

Promise only what you deliver

Learning local evidence for shading and reflectance

Matt Bell and William T. Freeman
Microsoft Electric Research Labs (MERL)
201 Broadway
Cambridge, MA 02139

Abstract

A fundamental, unsolved vision problem is to distinguish image intensity variations caused by surface normal variations from those caused by reflectance changes in a self-shading from point. A solution to this problem is necessary for machines to interpret images as people do and could have many applications.

We take a learning based approach. We generate a training set of synthetic images containing both shading and reflectance variations. We label the interpretations by indicating which coefficients in a stereable parametric representation of the image were caused by shading and which by point.

To analyze local image evidence for shading or reflectance, we study the output of two layers of filters, each followed by rectification. We fit a probability density model to the filter output using a mixture of factor analysis. The resulting model indicates the probability, based on local image evidence, that a given coefficient of any orientation and scale was caused by shading or by reflectance variations. We take the lighting direction to be that which generates the most shape like labeling.

The labeling allows us to reconstruct hand-drawn images containing only those parts of the input image caused

Intensity changes are due to surface normals contrast specular shapes when contour changes. Here, we restrict ourselves to shading from point.

Figure 1 (a) illustrates the problem: intensity changes are caused by the way which the point was viewed. Scene 4 effects. (b) shows the same location a after an attempt was made to reflect a over the rock. It is simple to see the task in the image (c), we want to develop a to do the same thing.

This problem has not yet been solved. Siskin and Adelson [1] solved the 2D world domain, based on heuristic rules, trees and contours, which were pre-learned other blocks world vision solutions, the analogous solution for real images.

Freeman and Viola [2] proposed a p shapes which generalized the elaboration of applied to explain images made by us. Their method assumed each image was or all main and couldn't receive an line

Some negatives

- Related prior work that you don't seem to be aware of.
 - "someone else did PCA on motion capture data before".
 - Better that you bring it up than the reviewers.

Quick checks you can do

- Does it deliver what it promises?
- Does it reference previous work in field?
- (note logical fallacy of rejection based on those faults).

What names should be on it, in what order?

- The people who contributed to the paper.
- Should your advisor's name be on it?
- What is a contribution?

- My rule of thumb: All that matters is how good the paper is. If more authors make the paper better, add more authors. If someone feels they should be an author, and you trust them and you're on the fence, add them.

Title?



Our title

- Was:
 - Shiftable Multiscale Transforms.
- Should have been:
 - Shiftable Multiscale Transforms, or, What's Wrong with Orthonormal Wavelets?

Author list

- It's better to be second author on a great paper than first author on a mediocre paper.
- The benefit of a paper to you is a very non-linear function of its quality:
 - A mediocre paper is worth nothing.
 - Only really good papers are worth anything.

Author order

- Some communities use alphabetical order (physics, math).
- For some it's like bidding in bridge.
- Engineering seems to be: in descending order of contribution.
- Should the advisor be on the paper?
 - Did they frame the problem?
 - Do they know anything about the paper?
 - Do they need their name to appear on the papers for continued grant support?

NIPS title word statistics

- For banquet talk, analyze words in title for ability to predict papers chance of acceptance.
- Most predictive of acceptance:
 - Bayesian, Gaussian.
- Most predictive of rejection:
 - Neural, network.

Ted Adelson on writing papers.

- (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.

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

Show simple toy examples to let people get the main idea

From
"Shiftable
multiscale
transforms"

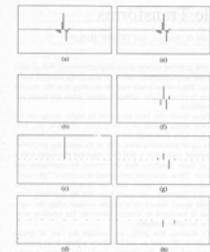


Fig. 1. Effect of modulation on the wavelet representation of an input signal, which is equal to one of the wavelet basis functions. (a) (b) Decomposition of the signal into three wavelet subbands. Plotted are the coefficients of each subband. Dots correspond to zero-value coefficients. (c) Same type signal, modulated one sample to the right. (d) (e) Decomposition of the shifted signal into three wavelet subbands. Note the drastic change in the coefficients of the wavelets, both within and between subbands.

Be kind and gracious

- My initial comments.
- My advisor's comments to me



Efros's comments

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 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.

342

Develop a reputation for being clear and reliable

- There are perceived pressures to over-sell, hide drawbacks, and disparage others' work.
- "because the author was Fleet, I knew I could trust it." [recent conference chair discussing some of the reasons behind a best paper prize].

Be honest, scrupulously honest

Convey the right impression of performance

Knuth

12. Motivate the reader for what follows. In the example of §2, Lemma 1 is motivated by the fact that its converse is true. Definition 1 is motivated only by decree; this is somewhat riskier.

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?

When describing the work of other people it is sometimes safe to provide motivation by simply stating that it is "interesting" or "remarkable"; but it is best to let the results speak for themselves or to give reasons why the things seem interesting or remarkable.

When describing your own work, be humble and don't use superlatives of praise, either explicitly or implicitly, even if you are enthusiastic.

Knuth

13. 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.

Knuth

24. The opening paragraph should be your best paragraph, and its first sentence should be your best sentence. If a paper starts badly, the reader will wince and be resigned to a difficult job of fighting with your prose. Conversely, if the beginning flows smoothly, the reader will be hooked and won't notice occasional lapses in the later parts.

Probably the worst way to start is with a sentence of the form "An x is y ." For example,

Bad: An important method for internal sorting is quicksort.

Good: Quicksort is an important method for internal sorting, because ...

Bad: A commonly used data structure is the priority queue.

Good: Priority queues are significant components of the data structures needed for many different applications.

Mermin

rule in your original manuscript.

Rule 2 (Good Samaritan rule). A Good Samaritan is compassionate and helpful to one in distress, and there is nothing more distressing than having to hunt your way back in a manuscript in search of Eq. (2.47) not because your subsequent progress requires you to inspect it in detail, but merely to find out what it is *about* so you may know the principles that go into the construction of Eq. (7.38).

The Good Samaritan rule says: *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 \mathbf{E} and the Lindhard form (3.51) of the dielectric function ϵ into the constitutive equation (5.13) ..."

The elements of style, Stunk and White



The screenshot shows the Bartleby.com website interface. At the top, there is a navigation bar with links for Home, Search, New, About, Reference, History, Sources, Quizzes, and English Page. Below the navigation bar, the breadcrumb trail reads: Reference > Usage > William Strunk, Jr. > The Elements of Style. The main content area features a small image of the book cover for 'The Elements of Style' by William Strunk, Jr. To the right of the image, the title 'The Elements of Style' is displayed in a large, bold font, followed by the author's name 'William Strunk, Jr.'. Below the title and author name, there is a short description of the book, stating that it is a classic reference book for students and conscientious writers. At the bottom of the page, there is a search bar and a 'CONTENTS' section with a link to the 'Bibliographic Record'.

<http://www.bartleby.com/141/>