

Writing Research Papers

Aaron Hertzmann

Why should you care about writing?

- You must communicate your work to the world
 - If people don't know about it, they won't use it
 - Readers have finite bandwidth
 - Increase probability of adoption
- Clear writing requires clear thinking; muddled writing is a sign of muddled thought
- For its own rewards
 - good writing can be a joy; bad writing is agony
 - Be kind to your readers
- These skills transfer to other endeavors

Goals for your paper

1. **Describe** ideas and experiments accurately and honestly
2. **Persuade** people of the approach

Writing is a lifetime skill

The first step

Read a lot

Know Your Audience

- Potential audience includes:
 - Experts in your area; people outside your area
 - Expert professors; advanced undergrads
 - Industry practitioners
 - Researchers now; researchers 10 years from now
- Title/abstract often used to decide whether to read the paper (and when searching)
- Audience knows a lot less about this than you
- The clearer and more self-contained the paper, the wider the audience

Know Your Audience

- Put yourself in the reader's shoes
- Imagine other papers you've read and how easy/difficult they were at first

Editing

In research, everything requires iteration,
iteration, iteration

The Importance of Editing

“I'm not a very good writer, but I'm an excellent rewriter.”

– *James Michener*

“The beautiful part of writing is that you don't have to get it right the first time, unlike, say, a brain surgeon.”

– *Robert Cormier*

“The time to begin writing an article is when you have finished it to your satisfaction. By that time you begin to clearly and logically perceive what it is you really want to say.”

– *Mark Twain*

“Proofreading is an art and a craft.”

– *The Chicago Manual of Style*, 14th ed.

Editing

- 90% of writing is editing
- *Delete every unnecessary word*
- Break down complex sentences
- Refactor sentences for clarity and flow
- Convert passive into active voice
- “Vermont ~~is a state that~~ attracts visitors because of its winter sports.”

Passive vs. Active

- **Passive:** “The algorithm is poor at sorting.”
 - Anything with is, was, will be, has been, etc.
- **Active:** “This algorithm performs poorly at sorting.”
- Active sentences are clearer and more direct; passive are more indirect.
- Avoid passive writing when possible without sacrificing clarity
- “We” is acceptable as a way of avoiding passive voice
 - ... but don't do this: “We then sort the vertices by height...” (unless you are manually doing the sorting yourself, and not the algorithm)

Parts of a paper

Inverted Pyramid of Journalism



Tell a Story

- Humans communicate through storytelling
- We are fascinated by stories
- Pose the problem, ask a question, pose a solution, note problems that arise, address them, denouement

The Three-Act Structure

Act 1: Introduce the hero, and the conflict

Act 2: Action is taken, building to an inevitable showdown

Act 3: Climax and resolution

e.g.: Syd Field, *Screenwriting*, Dell 1984

Title

- Summarize the paper in 2-10 words
- As short as possible, and no shorter
- Goal: encourage reader to read the paper
- Examples for discussion:
 - Video Textures
 - An Image Synthesizer
 - Learning Physics-Based Motion Style with Nonlinear Inverse Optimization
 - Removing Camera Shake from A Single Photograph
(first draft title was something like: “Blind Image Deconvolution By Multiscale Variational Search”)

Abstract

- Summarize the paper in a paragraph or two
- Include: contributions, approach, results, advantages.
- As short as possible, and no shorter
- Goal: encourage reader to read the paper
- First sentence: summarize the paper
 - Don't make me read an essay before saying what you do.
 - “In this paper, we introduce a new algorithm for computing the bounding box of a penguin.”
- Rest of paper should stand alone without abstract; repeating text is OK.

Introduction

Goal: provide context and encourage reader to read the paper

The introduction has several parts:

1. Background and motivation (1 paragraph)
2. Overview of the paper and contributions (1-2 paragraphs)
3. More details and summary of the approach
4. Summary of the results and conclusions

Intro: Background and Motivation

- Place your work in a larger context. Why is it worthwhile? Why isn't the problem already solved?
 - “Obtaining accurate reflectance properties of real-world objects from single-view video is an open problem with many applications in special effects and molecular gastronomy. Unfortunately, current methods are restricted to the diffuse and vegetarian case...”
- Sometimes it would take too long to provide a grand context, so don't bother
 - “A major open problem in computational geometry is computing the bounding box of a penguin.”

Intro: Do not do this

“The problem of X has attracted considerable interest in the research community.”

(is this sociology?)

“Computer graphics has made great strides in photorealistic rendering. However, an alternative approach has emerged, called *Non-Photorealistic Rendering* (NPR)...”

(don't just copy the same paragraph that appears in 100 NPR papers by now... say something relevant to the paper)

Intro: Do not do this

“This paper is organized as follows. In Section 2, we introduce our algorithm. In Section 3, we test our algorithm. In Section 4, we run some experiments. In Section 5, we cut and paste text from the Introduction.”

Does anyone actually benefit from this crap?

“I try to leave out the parts that people skip.” *–Elmore Leonard*

Intro: Managing expectations

Sell your work, but do not lie

- All claims should be supported with experiments, proofs, etc.
- Avoid grandiosity (“we propose a framework...”)
- Unfulfilled promises are the kiss of death

State your assumptions!

We assume that all image blur can be described as a single convolution; i.e., there is no significant parallax, any image-plane rotation of the camera is small, and no parts of the scene are moving relative to one another during the exposure. Our approach currently requires a small amount of user input.

Intro: Providing context

- Make the context and contribution clear
- The reader should know precisely why the paper is worth reading
(e.g., a paper with a geometry component and an HCI component might have novelty in only one of the two areas; make it clear what is novel)

Related Work

Several goals for this section:

- Acknowledge your debt
- Explain precisely how your work is different
- Stroke reader/reviewer egos
- Outline your perspective on the field (esp. if your paper is countering an orthodoxy)

“If I have seen further it is by standing on the shoulders of Giants.” –Isaac Newton

Related Work

DO

- Point out **both** advantages and disadvantages of related work
 - (provides context; defuses objections; is honest)
- Discuss all references that you cite

DO NOT

- Write a laundry list
- Bash the references
- Include irrelevant references
- Write a paragraph about a very peripheral work

Overview

Sometimes helpful if there are a lot of parts

3 Overview

The main idea of our work is to learn a probability distribution function (PDF) over character poses from motion data, and then use this to select new poses during IK. We represent each pose with a 42-dimensional vector \mathbf{q} , which consists of joint angles, and the position and orientation of the root of the kinematic chain. Our approach consists of the following steps:

Feature vectors. In order to provide meaningful features for IK, we convert each pose vector to a feature representation \mathbf{y} that represents the character pose and velocity in a local coordinate frame. Each motion capture pose \mathbf{q}_i has a corresponding *feature vector* \mathbf{y}_i , where i is an index over the training poses. These features include joint angles, velocity, and vertical orientation, and are described in detail in Section 4.

SGPLVM learning. We model the likelihood of motion capture poses using a novel model called a Scaled Gaussian Process Latent Variable Model (SGPLVM). Given the features $\{\mathbf{y}_i\}$ a set of motion capture poses, we learn the parameters of an SGPLVM, as described in Section 5. The SGPLVM defines a low-dimensional representation of the original data: every pose \mathbf{q}_i has a corresponding vector \mathbf{x}_i , usually in a 3-dimensional space. The low-dimensional space

Body

- *Teach* the reader
- Organize the paper with a logical flow
- Provide sufficient signposting to explain where you're going and to dive in
- Provide experiments and demonstrations to justify all of your main claims
- Compare with all relevant existing methods (and obvious, trivial extensions)

Models vs. Algorithms

- Whenever possible, describe abstractions before algorithms
 - Energy functions
 - Generative models
 - High-level goals
- Increases clarity and possibility for reuse
- Example: Lloyd's method

Redundancy

- Use both plain English and technical detail/intuition and details
 - Redundancy provides checksums (low-entropy codes are easily corrupted)
 - People have different styles of learning
- “Intuitively, variance is the average error over the data, and is computed as the average squared error: $\sigma^2 = \sum_i (y_i - f(x_i))^2 / N$ ”

Reproducibility

- Practitioners (e.g., skilled grad student) should be able to reproduce your work from the descriptions, down to the level of tuning parameters (if possible)
 - Don't assume some steps are obvious
- Release your code and data online (a delay is acceptable if you want to do follow-up work)
 - It doesn't have to be production-quality
- I know from experience: these things make a difference!

Figures

- Always appear *after* the first reference
- Reference in numerical order
- Provide visualizations of the model/results
- Make sure the text is legible (not 3pt, a common problem with MATLAB plots).
- Make figures that work in grayscale whenever possible

Discussion/Conclusions

- What can you say about the work that you couldn't before? What are the broader implications of the work?
- Don't just repeat the introduction/abstract
- If you cannot think of anything to say, just skip it (or keep it brief).

Future Work

I still haven't figured out what this is for

Possible goals:

1. Make clear what open questions remain in the domain of this work
2. Highlight the doors that your work opens up, create excitement
3. Make clear what limitations should be fixed
4. Give away free research ideas (preferably, ones you don't want)???
5. Denouement
6. Everyone has one, and you don't want to look weird

Discussion and Future Work

6 Discussion

We have introduced a method for removing camera shake effects from photographs. This problem appears highly underconstrained at first. However, we have shown that by applying natural image priors and advanced statistical techniques, plausible results can nonetheless be obtained. Such an approach may prove useful in other computational photography problems.

Most of our effort has focused on kernel estimation, and, visually, the kernels we estimate seem to match the image camera motion. The results of our method often contain artifacts; most prominently, ringing artifacts occur near saturated regions and regions of significant object motion. We suspect that these artifacts can be blamed primarily on the non-blind deconvolution step. We believe that there is significant room for improvement by applying modern statistical methods to the non-blind deconvolution problem.

There are a number of common photographic effects that we do not explicitly model, including saturation, object motion, and JPEG artifacts. Incorporating these factors into our model should improve robustness. Additionally, our method could be extended to make use of more advanced natural image statistics, such as correlations between color channels, or the fact that camera motion traces a continuous path (and thus arbitrary kernels are not possible). There is also room to improve the noise model in the algorithm; our current

Overcoming Writer's Block

- *Everyone* experiences it
- Vary your routine and find what works:
 - Imagine explaining it in person to someone
 - Write something awful and then edit it into shape
 - Find your productive times and/or spaces
 - Turn off/avoid distractions
 - Break it down into simpler tasks: write an outline first and then flesh out the details
 - Write on pen and paper
 - Bite the bullet
- WB is sometimes fear of failure
- Remember: **the first draft doesn't have to be perfect!**

References

- George Orwell, “Politics and the English Language” (1946)
 - <http://www.resort.com/~prime8/Orwell/patee.html>
- Strunk and White, Elements of Style
- Dupré, Bugs in Writing
- Chicago Manual of Style
- Kajiya, “How to Get Your SIGGRAPH paper rejected.” (1992)
 - <http://www.siggraph.org/publications/instructions/rejected>
- Shewchuk, “Three Sins of Authors in Computer Science and Math” (1997)
 - <http://www.cs.cmu.edu/~jrs/sins.html>