How to Write a SIGGRAPH Paper:
A Guide to Choosing a Good Research Topic, Doing the Research, and Writing It Up

David Salesin

December 2016

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).
SA '16 Courses, December 05-08, 2016, Macao
ACM 978-1-4503-4538-5/16/12.
http://dx.doi.org/10.1145/2988458.2988471

So, you want to write a SIGGRAPH paper?!

The bad news:

- It will be a ton of work!

The good news:

- You'll accomplish an incredible amount...
- And in a very short time!
- And you'll learn how productive you can be without sleep!

Increasing your odds

Only about 1 in 5 papers accepted.

What can be done to increase your chances?

LOTS:

- From choosing the right topic
  – even getting in the right *frame of mind* to do so –
- To writing each section
- To rules of grammar & diction

Generalizations & caveats

Q: Is all of this advice more broadly applicable?
A: I suspect so.

Q: Can you give me any guarantees?
A: Absolutely not!

Part 1:

Choosing a Good Research Topic

Choosing a research topic

THE most important decision you'll make

- and the greatest single indicator of success
- a *hugely* important decision!

Think about the time investment you're about to make

- nine months to two years (or more)

What if it's not an interesting problem to begin with?

So . . . how do you choose a good research problem?

How research is done

1. Someone comes up with a radically new idea...
2. Spawns follow-on work in this new wide-open field.
3. Gradually, the big fat easy problems get knocked off.
4. Remaining problems are more difficult — and smaller.
5. Inevitably, ten years later, there are still lots of people caught in these "research eddies" — *don't be one of them!*

The easiest way to get a SIGGRAPH paper

...and, more importantly, make a real impact:

Be the person with the brand-new way-out idea!

Since most papers are narrow improvements you may have gotten the opposite impression, but

- SIGGRAPH is *actively seeking* crazy new ideas.
- The "bar" is actually much lower for them.

Other advantages of opening new areas

1. You are a lot less likely to get scooped.
2. You've got a built-in head start on follow-up work.

Putting yourself in the right frame of mind

How do you come up with that brand-new really cool research project?

Not so easy:

- Cannot do it on the spot – more a frame of mind.
- Requires putting aside what you're working on now, and exposing yourself to new inputs & ideas.
- Read, go to seminars & talks, and – equally important – pursue your outside interests!

Putting yourself in the right frame of mind

Don't just experience these things passively:

actively engage by *asking questions!*

For instance, when reading a research paper, ask yourself:

- What are the limitations of the paper?
- How could they be overcome?
- Can ideas from this paper be combined with others?
- Does the paper provide a new "hammer"?
- Or do you have a hammer that could improve upon it?
- Is the paper one example of a more general process?
- If so, what other problems could it be applied to?

Putting yourself in the right frame of mind

May not be the best way to get work done on your immediate project.

No one paper or talk is likely to pay off with a brilliant idea.

But in the long this is the only way to stay creatively engaged:

Constantly filter the ideas around you.

Selecting the problem

If you've been doing all of this, then you probably have *lots* of research ideas.

Now, how to choose?

Here are 6 criteria to consider....

Selecting the problem

1. How important a problem is this?

   The more it would make a "big difference in the world," the better.

   If it's an incremental advance at best, I'm less interested.

Selecting the problem

2. How good an idea do I have for how to solve it?

The more I have a good sense for how to solve it, the better.

Of course, it can't be trivial either: then it's not research!
...though this is usually not an issue.

(If it's important and unsolved it's rarely trivial!)

Selecting the problem

3. How original is the problem?

The more original, the better.

If it's a well-trodden area, I'm less interested.

(Why would I think I can do better than a lot of smart people who have already been here?)

Selecting the problem

4. How well-defined is the problem?

The more I could put together a clear problem statement & convince the world I have indeed solved it, the better.

Selecting the problem

5. How much would the problem lead to interesting continuing future work?

   The more this will lead to other good problems, the better.

Selecting the problem

6. How interested am I in the problem, personally?

   Related: Do I have some "secret weapon" for solving the problem – something that makes me especially qualified?

Combining the criteria

The metric I use is the *product* of these different measures.

So if it fails even just *one*, then it may not be the best problem to pursue – at least *at this time*.

Try answering each of these questions ***in writing***!

Selecting the problem, cont.

Don't worry if the first problem you think of doesn't pass muster.

Whatever you work on, you will work on for a long time; spend time up front choosing the right problem to pursue.

Not unusual to spend 1/3 of total time on a research problem figuring out what the *right* problem should be!

Selecting the problem, cont.

Don't worry too much about *already* being an expert in your chosen research area — that's too limiting.

In whatever area you choose, you will have to spend time getting up to speed.

And there are benefits to being a "beginner": you may bring a fresh approach that's just what the area needs.

One other criterion

One other question I ask myself:

Why hasn't this problem been solved before?

I feel much happier if I can think of *some good reason, e.g.,*:

- Computers weren't fast enough until now
- No "hammer" like this had been available

Best if there's a good reason why no one else would even *have thought* to solve this problem before!

One other place to look...

Earlier we talked about the trend from seminal papers, to lots of good research exploration, to "research eddies."

Sometimes though a seminal work is so good — or appears to be so — that the research area becomes open-and-shut.

...for a while. Eventually, years later, there is some good follow-on work, and the research floodgates open.

So look for inspiration as well in classic, seminal papers — the ones that really impressed you — that also haven't received much follow up.

And one last thought...

Choice of problem is important, but won't determine your whole life's work – so don't obsess too much!

That said, it's one of the most important skills to develop:

Learning to choose a good problem is every bit as important as learning how to solve one, if not more.

Part 2:

Doing the Research

Research project stages

Like a game of chess, research can be divided into stages:

1. The opening game
   ...prepare your attack
2. The middle game
   ...pace yourself as you progress
3. The end game
   ...prepare a killer demo & write it up!

The opening game

What *not* to do first:

- Build a massive software framework
- Worry about efficiency
- Handle edge cases

Instead:

- Penetrate straight to the most *innovative* aspects & prove them out.
- Don't get stuck: may need to change the approach *...or problem!*

The opening game

Often the full-blown problem is too hard at first.

Find ways to simplify it, e.g.,:

- Use synthetic data
- Lower the dimension
- Find ways to run quick tests

...optimize the research's "inner loop"

The middle game

This is main part of your research, generally lasting months.

Be careful to pace yourself....

Alternate between implementing & reflecting:

Don't get stuck in just one or the other!

The middle game

When things are going well:

- Your research "has legs": it practically runs itself, dragging you along after it!
- Can become an extended state of "flow"
- Becomes a matter of prioritizing all your ideas for new directions to explore and results to achieve:
  Keep a written list so you can track & prioritize them!

The middle game

When things are *not* going so well:

- Not necessarily bad ... can be an opportunity to reflect while getting maintenance work done.
- Helpful to set up a schedule with intermediate deadlines.
- Remember: "Activity gives rudder"!

The end game

As the submission deadline approaches, it's time to:

- Produce those really killer results & demos!
- Get rid of hacks you won't want to have to write about!
- Prioritize the remaining research so that you can cover/claim as much ground as possible.
- Uncover additional applications.

The end game

Begin writing the paper early!

- Helps you sort out the big ideas from the hacks — and formalize the intuition behind any hacks that remain.
- Gives time to craft the "story" that you want to tell.

Keep a notebook!

Definitely worthwhile to write stuff up as you go along:

- Helps you keep track of all the difficult problems that had to be solved – useful for writing the paper.
- Helps you think ideas through.
- Helps you measure progress against earlier results.
- Helps you disseminate results & seek constant feedback from those around you.
- Helps you keep track of contributions to acknowledge.

Protecting your ideas

I prefer to err on the side of *sharing too much*, if anything:

- Allows you to get expert feedback early on.
- Helps you find out early if you're being scooped.
- Engenders a community spirit.
- Allows you to lay claim to the territory.

If you are scooped:

- Read through the paper calmly and critically.
- Chances are it's not as bad as it might first appear.

Managing your own emotions

It's a long road: there will be lots of ups and downs.

A hard problem is a good thing!

. . .but it's a very difficult thing to find a problem that's just hard enough.

Often, the problem will have to change if you're going to make any good progress on it — and this is OK.

Part 3:

Writing It Up

Why is writing well important?

Helps you clarify your own ideas.

If it's well written well, it can be appreciated by a larger audience:

...beyond your field

...and well into the future!

How will your work be judged?

...by your best paper?

...by your average paper?

...by your worst paper?

What's unique about a research paper

Writing a research paper is different from other technical writing you may have done:

It's not just about careful, clean exposition.

It's also about how you got there:

- *What the problem is* (potentially a major contribution)
- *How you arrived at the solution*: the deductive steps used

Why this is important

The discussion of how you arrived at the solution is important for several reasons:

1. Helps provide an *appreciation for the depth & difficulty* of the problem itself – especially if the problem is new.
2. Helps add *narrative* to the paper, making it interesting.
3. Provides a *framework* for discussing *alternate choices*.
4. Helps the reviewer evaluate the logic of your design & whether there is something *inevitable* about your solution.

Akin to giving the mean vs. the mean plus variance.

Brevity

All about simplicity: be as concise as possible!

Only introduce concepts (e.g., things with italicized names), notation, etc., when absolutely necessary.

Keep the notation simple as well!

Figures

Think of the figures as a *storyboard* for the whole paper: a separate track that stands on its own.

Readers should be able to read through just the figures first:

- They should all make perfect sense.
- They should comprise a summary of the entire paper.

Try presenting your work as a talk with just pictures before writing the paper, then reuse them for the paper's figures.

Mathematics

Very important: *All equations and pseudocode should be preceded or followed by an informal plain-language description of what they mean!*

The paper should make sense if you insert "grunt" for every equation or expression or pseudocode.

Think of the math as providing a *third* way of reading the paper if you want the most detail.

Title

One of the most important decisions you'll make!
Worth a good deal of time & thought—and multiple revisits.

You want something as short as possible that adequately describes the work & how it differs from previous work.

Title, cont.

For something essentially new, you can use something short:

Image Analogies

Escherization

Video Textures

Comic Chat

Title, cont.

For something that's a refinement of previous work, you'll need something longer:

> Environment Matting Extensions: Toward Higher Accuracy and Real-Time Capture
>
> Scale-Dependent Reproduction of Pen-and-Ink Illustrations

Title, cont.

Sometimes the length is there to describe what it's good for:

> The Virtual Cinematographer: A Paradigm for Automatic Real-Time Camera Control and Directing
>
> Hierarchical Image Caching for Accelerated Walkthroughs of Complex Environments

Title, cont.

Selection of a title helps focus the rest of the paper:

- Makes you realize what is truly *the most important* contribution of the work.
- Helps you structure the paper around it.

Authors

Can be tricky: sometimes hard to draw the line between an actual *co-author,* or someone more worthy of an acknowledgement.

My personal rules:

1. *A collaborator you have been meeting with regularly over an extended period of time should be a co-author. Chances are they've made some important contributions.*

Authors, cont.

2. *Someone who has implemented a portion of the code for the purposes of the project should be a co-author.*
3. *Someone you have consulted with a few times, or who offered feedback on the paper, is probably more worthy of an acknowledgement.*
4. *If in doubt, err on the side of inclusion.*
5. *Put main contributor first, then order by seniority.*
6. *Don't stress out over author order too much!*

Affiliations

Sometimes tricky as well. My rules:

- *Again, err on the side of being generous.*
- *List* all *places in which you worked on the project.*
- *If the work was done at* A *but written up at* B*, list* both*.*
- *If the work was done entirely at* A *but now you happen to work at* B*, list* just A*.*

Abstract

Keep it terse; make every word count!
Go over it, again and again, with a fine-toothed comb.
Keep in mind: this is the one part of the paper that everyone will read, even those who don't look at the rest.
Write it after the rest is finished: much easier.
Keep it to a single paragraph — or at most two.
Not an introduction: don't recapitulate the whole set-up.
Get right to the point — then stop as soon as it's made!

Abstract, cont.

First sentence should be the one-line takeaway:

> *This paper describes a new framework for processing images by example, called "image analogies."*

Abstract, cont.

First sentence should be the one-line takeaway:

> *We investigate a new approach for reproducing color images.*

Abstract, cont.

First sentence should be the one-line takeaway:

> *This paper introduces and presents a solution to the 'Escherization' problem: given a closed figure in the plane, find a new closed figure that is similar to the original and tiles the plane.*

Abstract, cont.

Second sentence is generally a summary of the approach:

> *The framework involves two phases: a* design phase*, in which a pair of images, with one image purported to be a "filtered" version of the other, is presented as "training data"; and an* application phase*, in which the learned filter is applied to some new target image in order to create an "analogous" filtered result.*

Abstract, cont.

Second sentence is generally a summary of the approach:

> *Rather than mapping the colors in an image onto the gamut of colors that can be printed with cyan, magenta, yellow, and black inks, we choose the set of printing inks for the particular image being reproduced.*

Abstract, cont.

Be sure to (briefly) mention all significant contributions

...usually just requires an extra sentence or so.

Summarize any applications or results.

Optionally give a one-sentence take-away:

The result is a more compelling pen-and-ink illustration than was previously possible from 2D reference imagery.

That's it—stop!

(Optional) quotation

Occasionally, there's a really nice definition or quotation you can use to lead the reader in to the subject of the paper.

(Optional) quotation

Can serve a number of purposes:

1. It could be a statement by an authority to motivate the problem — or suggest the difficulty of the problem.

There is no landscape that we know as well as the human face.... Every detail of the nose, eyes, and mouth, every regularity in proportion, every variation from one individual to the next, are matters about which we are all authorities.

— *Gary Faigin [14], from* The Artist's Complete Guide to Facial Expression

(Optional) quotation

Can serve a number of purposes:

2. It could be a statement that you present with some irony: e.g., someone who says the problem is easy.

It is of interest ... that, regardless of the number of impressions, the inks may be selected solely on the basis of their color gamut. Their colors need not be cyan, magenta, and yellow; nor is it required that they be transparent. The way is therefore opened for entirely new printing processes.

—Hardy and Wurzburg, 1948 [6]

(Optional) quotation

Can serve a number of purposes:

3. It could be someone talking about something completely different, but which actually has some resonance:

A native talent for perceiving analogies is ... the leading fact in genius of every order.

—William James, 1890

Introduction

First sentence:

- Your one opportunity to capture the reader's imagination.
- Your one chance to be punchy or flowery.
- Purpose: Take the reader from the world of all possible ideas into the world of this paper:

Tilings are as old as civilization.

Introduction

First sentence:

- Your one opportunity to capture the reader's imagination.
- Your one chance to be punchy or flowery.
- Purpose: Take the reader from the world of all possible ideas into the world of this paper:

Watercolor is like no other medium.

Introduction

First sentence:

- Your one opportunity to capture the reader's imagination.
- Your one chance to be punchy or flowery.
- Purpose: Take the reader from the world of all possible ideas into the world of this paper:

> *Analogy is a basic reasoning process, one that we as humans employ quite commonly, and often unconsciously, to solve problems, provide explanations, and make predictions.*

Introduction

After 1-3 paragraphs, you should narrow in on *the problem*.

Try to state the problem definition both informally & formally:

Abstract

This paper introduces and presents a solution to the "Escherization" problem: given a closed figure in the plane, find a new closed figure that is similar to the original and tiles the plane....

Introduction

... **Problem** ("ESCHERIZATION"): Given a closed plane figure S (the "goal shape"), find a new closed figure T such that:

1. T is as close as possible to S, and
2. copies of T fit together to form a tiling of the plane.

Problem ("IMAGE ANALOGIES"): Given a pair of images A and A' (the *unfiltered* and *filtered source images*, respectively), along with some additional *unfiltered target image* B, synthesize a new *filtered target image* B' such that

$$A : A' :: B : B'$$

In other words, we want to find an "analogous" image B' that relates to B in "the same way" as A' relates to A.

Introduction

After the problem statement, be sure to cover:

- *Why the problem is difficult (...part of the research!)*
- *Why the problem is important*
- *Why previous approaches are insufficient*
- *Your approach to solving the problem – and not just the solution, but your reasoning behind it as well*
- *The basic contributions of the paper (...make your reviewers' job easy!)*
- *What the applications are – how you see it getting used*
- *Drawbacks & limitations (...important not to oversell!)*

Related work

Your mission: pull together all the related threads of work.

Can be exciting to juxtapose work from disparate fields.

Sort by "topic," then cover each topic chronologically.

(Topics could be the various subproblems, fields of work, or some other categorization.)

Put each topic in a separate paragraph; don't worry if lengths are uneven.

Describe how each work is related to yours.

Err on the generous side. And don't fall into the trap of criticizing earlier work that you are building upon!

Overview

Skip it – unless truly warranted!

Sometimes it could just be a paragraph at the end of the Introduction. (Or not even that.)

If you're going to have one, make it informative.

Background

Helpful to have this separate section to lay the groundwork for understanding the technical contributions.

...also helps make clear what's actually new!

...and can be a nice contribution in and of itself.

Write this section *lazily:* should have just the info necessary to make the paper standalone.

Think about developing it backwards – *e.g.,:*

"isohedral tilings" → *"isohedral"* + *"tilings"*

"isohedral" → *"transitivity class"* → *"symmetry"*

Algorithm

Before jumping into your solution, use narrative to lead the reader from "Background" to "Algorithm."

e.g., from [*Computer-Generated Pen-and-Ink Illustration*]:

- Discussed how one might implement the "Background" 's various principles as part of an automatic system.
- Identified key differences from photorealistic rendering ("dual nature of strokes," "need to combine 2D & 3D").
- Discussed which parts of traditional rendering pipeline could be left as is, and which would need to change.

Algorithm

This high-level analysis can be a contribution.

... i.e., the writing itself can be a contribution!

Could be done as a complete section, or as a way to introduce each part of the algorithm as you go along.

Algorithm

Generally best to organize by sub-problem, one per section.

Generally cleanest to use pseudocode, along with plenty of plain-language text, including:

- Inputs
- Outputs
- High-level description of how the pseudocode works
- Boundary cases & additional details

Results

Long before writing the paper, think hard about the cleverest ways to show off your research.

Often, a result that's just out of reach will motivate new, interesting research.

Be sure to include results demonstrating each of your claims.

Extra credit for showing something "unexpected."

Caption your results so that they stand on their own; the paper's text should provide an extra level of detail.

Don't forget to include running time, memory, and how performance scales.

Summary/Discussion

This section is a chance to reflect on research accomplished, while pointing out limitations & work still to be done.

Don't provide a summary unless you have something *new* to say: *...the mark of a good summary is revelation!*

May be a good chance to point out additional applications.

...keep it modest though, and don't overstate!

Provides a chance to compare/contrast to other work, now that the full context has been revealed.

Future work

Can be woven in with Discussion, or set off separately.

separate paragraphs, with 2-3 word headers, works well.

Generally, I order future work from most immediate enhancements, to ideas for furthest-out research.

True, you're giving away ideas, but also laying claim to the territory.

Acknowledgements

Be terse but generous.
Include everyone who has helped in some significant way.
Easier if you keep a running list while doing the research!

References

Avoid abbreviations:

...*spell out whole names of journals*

Give as much information as possible:

...*ISBNs, etc.*

Strive for consistency

Appendix

A good place to put details *essential for re-implementing* the work that are unnecessary for understanding the paper.

Even here, keep it terse, but clear.

...don't ignore any of the rules on writing!

Notes on mathematical writing

Symbols in different formulas should be separated by words:

Bad: Consider S_q, $q < p$

Good: Consider S_q, where $q < p$

Notes on mathematical writing

Never start a sentence (or even a phrase—i.e., anything after a comma) with a symbol or mathematical expression:

Bad: x represents...

Good: The value x represents...

Notes on mathematical writing

All variables/notation *must be defined* when it is first used!

Don't get carried away with subscripts:

> ...spend time designing notation that is as simple and sensible as possible!

Spell out numbers when used as adjectives but not as nouns:
Good: "There are two approaches to this problem."
Good: "In each iteration, the variable x is incremented by 2."

Notes on mathematical writing

Don't use colons before displayed equations unless the grammar of the sentence calls for them.

Good: Can be written as follows:

Bad: Let:

Other writing issues

Avoid the use of “this” on its own.

Bad: This is the reason that...
Good: This theorem is the reason that...

Other writing issues

Enclose parenthetic expressions between commas.

Bad: My wife Sylvia...
Good: My wife, Sylvia,...
Good: My brother Jeremy ... (I have two brothers.)

Other writing issues

"That" is the defining pronoun; "which" is the non-defining pronoun.

Good: The iPhone that is broken is on the table.
Good: The iPhone, which is broken, is on the table.
Bad: The iPhone which is broken is on the table.

Avoid the "wicked which."
In general, if you can replace "which" with "that" without changing the meaning of the sentence, you should.

Other writing issues

"Fewer" is for countable things; "less" for uncountable things.

Good: I would like less milk.
Good: I have fewer students than last year.
Bad: I have less students than last year.

Other writing issues

Don't use jargon unnecessarily—even when writing for people in a field who would understand it.

Don't praise your own work:

The following algorithm provides a very elegant solution

Footnotes

See first whether the idea could be worked into the text

Make sure text does not depend on footnote (or parentheticals)

Don't slide essential qualifications into a footnote:

...footnotes must only enlarge, never mitigate!

Reserve the footnotes for just a single type of reader:

e.g., don't mix up footnotes for the expert with footnotes for the novice

Capitalization

Don't overuse.

Capitalize all important words in the title.

Capitalize "Section 1," "Figure 2," etc.

But just capitalize the initial letter of section headings.

Finale*

Finale*

* *"How Everything You Need to Know About Research is in* The Wizard of Oz"

Rule #1

In choosing a research area, follow your HEART

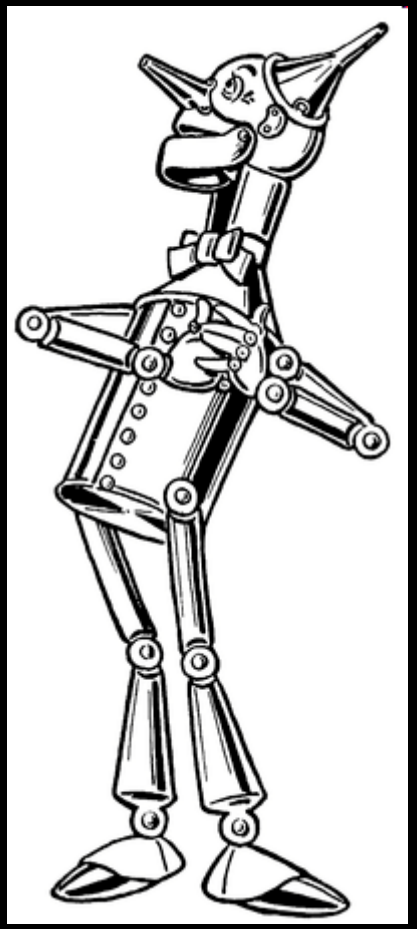

Rule #2

Have the COURAGE to leave the mainstream

Rule #3

Pose your research problem WISELY

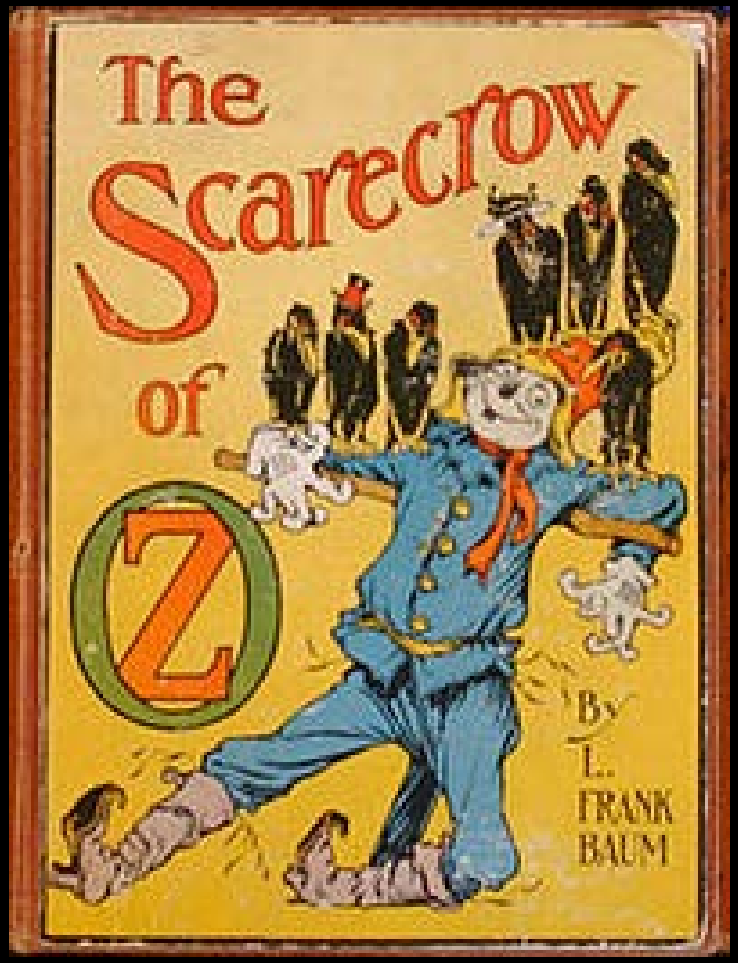

Rule #4

Solve hard problems with the best tools you have, strong collaborators, and above all PERSEVERANCE.

Rule #5

Embrace hard problems

Bibliography

Knuth, Donald E., Tracy Larrabee, and Paul M. Roberts. *Mathematical Writing*. 1989. ISBN 0-88385-063-X.

Strunk, William Jr. and E.B. White, *The Elements of Style, fourth edition*. 1999. ISBN 0-205-30902-X.

van Leunen, Mary-Claire. *A Handbook for Scholars, revised edition*. 1992. ISBN 0-19-506954-4.