

Guiding Principles for Communicating/Participating in Science

ORIGINAL SLIDES: NICOLE IMMORLICA, MICROSOFT

PRESENTER: RAF FRONGILLO, CU BOULDER

Before we begin...

1. I edited the slides a lot; credit→Nicole, blame→Raf
2. The advice I'll present is subjective; others could and do disagree
3. Most content is really about science more broadly, not just learning theory

Recurring Theme: focus on **why**.

understand the **generative process** behind good talks/websites/papers/etc

outline

(Some) aspects of communication/participation:

1. Papers you write
2. Reviews you write
3. Talks/Posters you present (if time)
4. Your website (skip)
5. In-person networking (skip)

Overview guiding principles/sample advice for each.

1. Writing Principles

“Writing is thinking. To write well is to think clearly. That's why it's so hard.” --David McCullough

Main goals:

- Communicate your results and why they are awesome
- Engage the reader

Everything else is secondary:

- No specific template (although norms)
- Some specific advice follows...
- **...but these are not hard/fast rules. Ask why!**

typical paper anatomy

1. **Title**: phrase indicating what you did
2. **Abstract**: concise description of what you did
3. **Introduction**: motivate and informally define problem, emphasize intuition and key contributions
4. **Related Work**: how your work fits with the literature
5. **Model**: the setting / methodology you consider
6. **Results**: statements, important/interesting details
6. **Appendix**: remaining proofs, minor extensions

title

Purpose: Briefly indicate why someone might be interested in reading your paper.

title

Good

Bad

Descriptive (but brief)

Vague

title

Bad:

Information Aggregation in Social Networks, Feldman, Immorlica, Lucier and Weinberg, 2014, working paper.

Good:

Reaching Consensus via non-Bayesian Asynchronous Learning in Social Networks, Feldman, Immorlica, Lucier and Weinberg, APPROX 2014.

title

Purpose: Briefly indicate why someone might be interested in reading your paper.

Example: Is it OK to be catchy/funny?

- Sometimes--if it is also descriptive.
- *Of the People: Voting is more Effective with Representative Candidates*, Cheng, Dughmi, and Kempe, EC 2017.
- *Prophet Inequalities made Easy: Stochastic Optimization by Pricing non-Stochastic Inputs*, Dutting, Feldman, Kesselheim, Lucier, FOCS 2017.

abstract

Purpose: Help someone understand what's in the paper (normally targeted at an expert).

Questions people will try to answer from your abstract:

- Am I qualified / do I want to review this paper? (!)
- Is this paper likely relevant to my research?
- Does this paper sound interesting?

abstract

Good	Bad
concise	wordy
factual/accurate	overstating/grandiose
identifies keywords	no (relevant) terminology

Pet Peeve: Abstracts that are really introductions.

abstract (bad)

In a social learning setting, **members of a society share their experiences to help others make better choices.**

Following the established path can boost an individual's utility but it can hurt the society as a whole since

other options of higher value may never be explored.

We show that when the population is diverse, this issue can be avoided as people may not be satisfied with the available choices and look for alternatives.

High diversity, though, comes at a cost as past experiences become less valuable. ----->

abstract (bad) (ctd.)

We model these situations in a standard setting of consumer search introduced by Weitzman and study how different diversity levels compare with each other. We ... and **quantify how the socially optimal diversity level changes** Moreover, while high diversity can lead to **anarchy and confusion** in typical situations, we show that it can be **really beneficial** in settings where society may accidentally uncover a unanimously accepted **hidden gem**.

abstract (good)

We introduce a general model of bandit problems in which the expected payout of an arm is an increasing concave function of the time since it was last played. We first develop a PTAS for the underlying optimization problem of determining a reward-maximizing sequence of arm pulls. We then show how to use this PTAS in a learning setting to obtain sublinear regret.

abstract

Purpose: Help someone understand what's in the paper (normally targeted at an expert).

Example: Should I sell the main results?

- You could state objectively why the main result is interesting, so an expert knows what's the point.
- You shouldn't go overboard, you have an entire introduction for that.
- E.g. "This is the first constant-factor approximation."

introduction

Purpose:

- Set the stage / motivate the general subarea
- Motivate the specific problem (the “gap”)
- State contributions and **why they solve the problem**
- Other context for why the paper is cool

Can be tricky to balance all of this!

Here's something that will help...

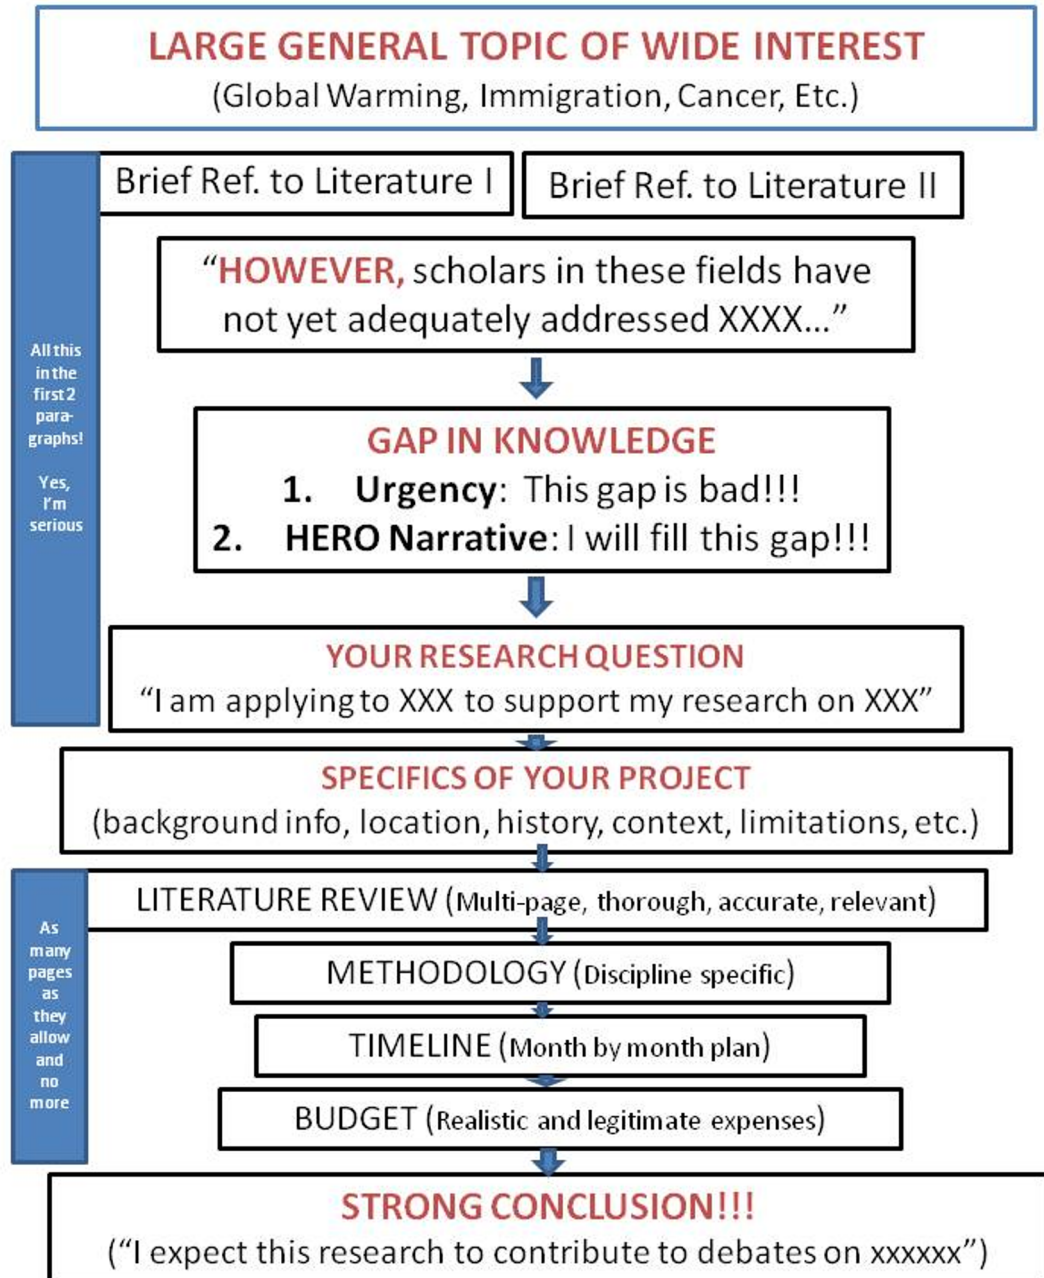
"The Foolproof Research Proposal Template"

introduction

Must-read:

Dr. Karen's
Grant Template

(google it)



introduction

Good

motivation from
practice
or existing literature

you fill an important gap

identify take-aways
and key intuition

Bad

flimsy/cartoonish, or worse,
NO MOTIVATION??

you do some stuff

overly-precise statement of
results and techniques

Pet Peeve: Laundry lists of results with no motivation.

introduction

Purpose:

- Set the stage / motivate the general area
- Motivate the specific problem (the “gap”)
- State contributions and why they solve the problem
- Other context for why the paper is cool

Example: Should I sell the main results?

- Absolutely! Don't be afraid to tell the reader exactly why it's cool.

introduction

Purpose:

- Set the stage / motivate the general area
- Motivate the specific problem (the “gap”)
- State contributions and why they solve the problem
- Other context for why the paper is cool

Example: Should I overview techniques?

- **Only do so to emphasize a why.**
- Ex: “We first use a reduction of XYZ, then Chernoff bounds, then some calculus” contributes nothing.
- “The key to our approach is the recent reduction of XYZ, developed for an unrelated problem.”

introduction

Purpose:

- Set the stage / motivate the general area
- Motivate the specific problem (the “gap”)
- State contributions and why they solve the problem
- Other context for why the paper is cool

Example: Should I overview techniques?

- Most theory papers choose to do this. Not required.
- Helpful if there is something exciting/digestible/etc.
- Really not helpful if the reviewer can't understand it.
 - Always ask: **what would a reader get from this?**

Related work

Purpose: Provide context for your work

- Most related stuff (ideally) already covered in intro
- Also to assign scientific credit for prior work.

NB: people you cite may be matched to review your paper -- use this to be inclusive, not exclusive

related work

Good

Bad

comprehensive

skimpy

describes connections
to this paper

reads like disjointed list
of abstracts

cites work from
multiple fields

unaware of related
literature

Related work

Purpose: Provide context for your work

- Most related stuff (ideally) already covered in intro
- Also to assign scientific credit for prior work

Example: How much detail should I give?

- Enough to make your point!
- Ex: “Cai and Daskalakis give a PTAS for a single unit-demand buyer with independent MHR item values, to the optimal deterministic item pricing.”
- Useful if you give a PTAS for a related problem.
- **Not** useful just because you study pricing.

Related work

Purpose: Provide context for your work

- Most related stuff (ideally) already covered in intro
- Also to assign scientific credit for prior work

Example: How much detail should I give?

- Enough to make your point!
- Ex: “Works such as [CaiD11, ...] also provide approximations in different models to ours.”
- Useful if you study unrelated pricing problem.
- **Not** enough if reviewer might reasonably wonder what your work contributes over CaiD11.

Model

Purpose: Start being formal

- Most intuition (ideally) already given in intro
- Need to be precise, but also clear.

NB: No results / theorem statements here!

Model

Purpose: Start being formal

- Most intuition (ideally) already given in intro
- Need to be precise, but also clear

Example: Should I give an example? If it serves a purpose.

- “The buyer’s utility is $v - p$. So for example, if $v = 5$ and $p = 1$, the buyer’s utility is 4.” **Useless!**
- “The buyer’s utility is $f(v) - g(p)$, for f, g convex. Observe that our model captures quasi-linear utilities ($v-p$) when $f(v) = v$ and $g(p) = p$. We will use this as a running example to illustrate the main ideas.” **Helpful!**

results

Good

intuition in main text

interesting proofs
that build intuition

illustrative examples

Bad

list of theorems/proofs

boring proofs included
because they're short

long unintuitive proofs

Results/Proofs

Ask yourself: Do I **want** the reviewer to read this?

Guiding Principle: **The entire body should be engaging.**

If you're bored writing, reviewer will be bored reading.

Make it exciting! Or maybe it belongs in appendix.

Try to make body as effortless to follow as possible.

If ideas too complex, distill main digestible aspects.

Appendix

Purpose: Verify omitted details.

- Yes, few people will read, but... details matter.

Style suggestions:

- Easy to read (not nec beautifully written)
- Organized

Pet peeve: appendices with typos, or worse, serious errors, which were obviously never proofread.

appendix

Good

clean to follow, even if
not engaging

pointers to/from body

Bad

unreadable

disorganized

2. Reviewing Principles

purpose

For conferences, your job is **not** to directly decide whether to accept/reject the paper.

Your job is to

1. give arguments/evidence/information to the PC so they can decide whether to accept/reject
2. give feedback to the authors (for science!)

Brief interlude from Raf's colleague Dan Larremore
(emphasis: give **constructive** feedback!)

typical content of a review

1. Summary
2. Pros/Cons
3. Evaluation
4. Recommendation

1. Summary

Describe paper in 1 paragraph. Enough context to explain **why authors think its exciting**.

Thought experiment: Keep it factual. Would the authors agree? (If so, it's a sign of a serious misunderstanding.)

Excellent: So clear that PC doesn't need re-read intro.

Bad: Too short (why bother?) or mixed with your opinion

2. Pros/Cons

Still factual (but sometimes blends into eval)

What are upsides or limitations? Is there an error?

Excellent: “I might be misunderstanding something, but as stated, it seems that Theorem 2 is false. Here is a sketch of a counterexample. Is it possible that the authors meant to place additional assumptions?”

Point: You may be about to kill the paper. If it needs to be done, it needs to be done. But **be thoughtful.**

2. Pros/Cons

Still factual (but sometimes blends into eval)

What are upsides or limitations? Is there an error?

Bad: “Theorem 2 is false, integral might diverge.”

(Often resolved by: “It is easy to verify that if integral diverges, results still hold with notational updates.”).

Point: Totally valid minor concern to be rigorous with divergence (and this should be raised). But don't kill papers for oversights which can be easily resolved.

3.

Evaluation

Now for your **opinion**.

Quality of **results**, general interest in **techniques**, quality of **presentation**.

Results: (In your opinion)

- What makes the results significant (or not)?
- Any context the PC needs to appreciate?

3.

Evaluation

Now for your **opinion**.

Quality of **results**, general interest in **techniques**, quality of **presentation**.

Techniques: (In your opinion)

- Will they help you/others solve problems?
- Did you find them engaging/illuminating?
- **Avoid:** “How much hard work?”
- **Simple results can be even more impactful;**
Nash’s equilibrium paper was 2 pages

3.

Evaluation

Now for your **opinion**.

Quality of **results**, general interest in **techniques**, quality of **presentation**.

Presentation: (In your opinion)

- Did the body back up the sale of the intro?
- Do you understand what you/the authors want to?
- A pleasure to read?
Okay to give up / tank a paper if truly awful to read

typical content of a review

1. **Summary**. Context for why authors find it exciting
2. **Pros/Cons**. Contributions, errors
3. **Evaluation**. Quality of: results, techniques, writing
4. **Recommendation**. Accept/reject? **Why or why not?**

by the way...

You will be **judged** by the quality of your review

People who read your review: Authors, other reviewers, PC members; increasingly, the world...

For some PC members, this may be the first time they see your name. Make a good first impression!

3. Presentation Principles

(and poster principles)

Let's hear from Dan Larremore...

https://www.cs.princeton.edu/~smattw/AMW20/Workshop_GivingATalk.pdf

Original slides left here for reference

4. Website Principles

Why should you make a website?

- Facilitate connections!
Say you give a good talk/poster presentation and Daniela notices. She wants to learn your year, advisor, other work you do. Make it feasible!
- Good to match norms.
In CS, senior PhD students usually have websites. Might look odd or bad if you don't.

James R. Wright

Hello, I'm James Wright. I am a [postdoctoral researcher](#) at [Microsoft Research in New York City](#). In July 2018, I will start as an Assistant Professor at the [University of Alberta](#). I completed my Ph.D. at [UBC](#) in 2016, advised by [Kevin Leyton-Brown](#).

Research

My research interest is in using data-driven methods to understand what is, behavior in interactions with other participants. My long-term goal is designing algorithms for media agents rather than idealized, perfectly rational game theoretic agents.

1. Picture!

2. Bio/Affiliation

Curriculum Vitae

My [academic CV](#) is available as both an [HTML page](#) and a [PDF document](#). I also have a public [Google Scholar citations page](#).

Publications

1. [Predicting Human Behavior in Unrepeated, Simultaneous-Move Games](#).
James R. Wright and Kevin Leyton-Brown.
Games and Economic Behavior, Volume 106, pages 16–37, November 2017.
(supersedes Wright & Leyton-Brown [2010, 2012])
2. [Learning in the Repeated Secretary Problem](#).
Daniel G. Goldstein, R. Preston McAfee, Siddarth Suri, and James R. Wright.
ACM Conference on Economics and Computation (ACM-EC), 2017.
3. [Deep Learning for Predicting Human Strategic Behavior](#).
Jason Hartford, James R. Wright, and Kevin Leyton-Brown.
NIPS 2016.

Oral presentation

4. [Incentivizing Better Behavior](#).
Worse.
Xi Alice
*Workshop on
Economics and Computation*, 2016.

4. Research/Publications



Contact

james.wright@ualberta.ca

3. Contact

4. Website Principles

Purpose: Introduce yourself. Use this to guide design.

Example: What if I don't have any publications yet?

- OK to omit. OK to put “coming soon!”
- People still want to know you without publications.

Example: Should I list unfinished manuscripts?

- Reason: excited about not-yet-published work.
- Reason: highlight research activity.

4. Website Principles

Purpose: Introduce yourself. Use this to guide design.

Example: Should I include personal information?

- Depends on what you want others to see!

Example: What if I'm bad at HTML?

- Ask to copy a friend's source code
- Use generic (tidy!) layouts, or bare bones

5. Networking Principles

Context: Researchers love meeting new students.

- But also super busy. Just be respectful of time.

Sample starters:

- Thanks for the great talk. I was wondering, how did you manage to prove/analyze X?
(Better yet: are you aware of the related work X?)
- I'm working in <your area>, and wondering if you're aware of any prior work which does XYZ?
- If you happen to bump into them, just say hi!

5. Networking Principles

Context: Researchers love meeting new students.

- But also super busy. Just be respectful of time.

Best Case: they want to hear about you. Be ready!

- Good to have an elevator pitch prepared.
- “So what do you work on?”
- “Tell me about one of your favorite projects.”

Appendix: How to Parse Feedback

Context: You'll constantly get feedback forever.
Collaborators, colleagues, reviewers, course evals, ...

Some feedback is amazing and easy to parse.

- “I think you should do XYZ because ABC.”
- And you completely agree immediately.

Most is not. Especially from reviewers.

who didn't attend this talk ;-)

- Vent. Go for a run. Then use it to improve!

feedback examples

Theme of advice: **Why** did this person give this feedback.

“I can’t follow the proof of Theorem 1, why does $Ax = b$?”

- If you clearly stated why $Ax = b$, OK to complain!
- But don’t **only** complain. Make it clearer.

“The results are fine, but incremental compared to [ABC].”

- If [ABC] is completely different, OK to complain!
- But **adjust** your next draft to better distinguish.

personal(?) anecdote

- Senior prof offered to skim Alice's job talk.
 - ~50 minutes on her thesis work.
- **Pretty blunt appraisal (paraphrased):** "So you solved these hard open problems. But your work is not too relevant or exciting and it's not clear what's next."
- **That hurt Alice a lot, but useful nonetheless:** others may get the same impression; really need to address.
- **Point:** Alice disagreed with the evaluation, but the feedback was still helpful. Radically improved her talk.

This has been: Guiding Principles for Communicating/Participating in Science

YMMV; remember: credit→Nicole, blame→Raf

Recurring Theme: focus on **why**.

understand the **generative process** behind good talks/websites/papers/etc

Thanks for listening / hope this was useful!