Research Paper Pitching Worksheet

1. Topic/Working Title:

Your title should capture both the problem that your solving as well as give a hint as to the methods employed. E.g. "Non-parametric methods for lens calibration".

_____

Conference/Journal: Page limit (min/max): Submission deadline:

_____ _____ _____

2. Who will work on the paper, and in what roles?

Who will contribute to this paper? Contributions include writing, use of other's software, hardware support (including robots), advising, and data collection. Who will be listed as a co-author (as opposed to an acknowledgement or citation?) Include UROPs, faculty advisors, etc.

Who	Contributions to paper	Co-Author?

3. What are the core contributions/claims of the paper?

Aim for three-ish. These can include conceptual insights, specific algorithms or mathematical discoveries, the creation of evaluation methods superior to previous ones, or competitive evaluations between multiple algorithms (i.e., a survey). In most cases, these will involve a "technological lever" (a clever idea, new algorithm, etc.) that explains why you are successful where others have failed. These claims/contributions are the reason for your paper's existence, and your job as an author is to make interesting/worthwhile claims and provide evidence to support them.

1.

2.

3.

4. Identify two related work papers/patents/theses, NOT from your lab:

(Do not go nuts doing a literature review. But take a quick look-around.)

5. How will you evaluate your method?

What experiments will you run? What methods (include strawmen and state-of-the-art) will you compare to your proposed method? What benchmark datasets will you use? What performance metrics? How will you present your data? (Tables bad, squiggly-line plots good.) Sketch anticipated plots (including what the axes are). Use additional pages if necessary.

6. What is the cool result that will excite the readers the most?

(What figure goes on page one? What "iconic" figure do you want your paper forever associated with?) SKETCH the figure.

7. What resources will you need to do this research?

Equipment:

Time (including yours):

8. What is the overall risk given the available time?

Low risk; much of the work is complete and just need to write.	Some risk due to either untested ideas or significant work remaining	Highly speculative and untried ideas/High Risk

0 1 2 3 4 5 6 7 8 9 10

9. What is the relationship between this project and your prelim/thesis/funded project?

A cartoon sketch of typical ICRA/IROS paper.

Title.
The title should capture the essential problem being solved and give some flavor of the approach. But use of jargon should be minimized-- the title will be read by people from a broad audience.

Page 1.
Abstract - Succinct answering of the four questions, no more than two (short) paragraphs.

Introduction - An elaboration of the four questions, put into context. You're laying the groundwork so that the reader understands your claims, and so that they'll understand the significance of your claims. This is not a literature review/previous work section, try to keep things a bit more abstract.

Claims - A list of 3ish claims/contributions of your paper. You want the reader to say "Well, gee, this sounds great. If they did all this, then the paper should definitely be published." (The Methodology and Evaluation sections are your primary weapons for supporting these claims.) I like to try to make these fit in the lower right of the paper.

Iconic Figure (in upper right) - Put a figure on the first page, even if it is a results figure that is "premature" with respect to the text flow. This is really a partner to your abstract and title in helping a prospective reader understand what your paper is about.

[On a six page paper, you should not need a "roadmap" where you say you will first talk about prior work, then the method, then show results, etc.]

Page 2.
Prior work - Describe related work. It is your goal to convince the reader that you are a scholar--- that you understand the state of the art. This section should be *about* the prior work, not about your proposed work. You can add value by organizing prior work into an ontology (perhaps there are two general types of solutions?). It is occasionally helpful to briefly identify how your work deviates from this previous work. ("A standard method for X is Y, which does A and B to solve the problem. Our proposed method also employs an A-like method, but uses C instead of B.")

You should generally be positive about earlier work (but not sycophantic). The existence of previous work similar to your own approach should not be viewed as a "weakening" of your paper, but as evidence that the problem and general approach are important.

Page 3-4.
Describe your problem formulation--- make it very clear exactly what problem you're solving, including cost functions, probabilistic formulations, etc. Along the way, help the reader understand your notation and any other important conventions.

Include algorithm blocks if appropriate. Include figures showing intermediate processing steps.

Your goal is for every reader to be able to reproduce your algorithm. That said, it's often impossible to address every detail, and getting the high-level ideas across is critical--- even if it means that you gloss over a corner case. (But make sure that the full method, including wrinkles and warts, is at least available for download.)

Page 4-5.
Evaluation. You are presenting evidence to support your claims. Provide comparisons against state of the art methods AND "an obvious thing to do". The latter can be a strawman, but may be easier to understand to the average reader (and be useful for intuition building) than a state of the art method. E.g., if doing a classifier, throw in a nearest neighbor method and a "random guessing" classifier.

Page 6.
Conclusion. Your job here is to briefly draw the reader's attention to the fact that you've made some claims about an interesting/important problem, and that you've supported those claims with evidence. Keep it brief... no more than two or three column inches.

No future work. Future work sections are lame. Anyone can speculate as to work they *might* do in the future. If there are weaknesses in your algorithm, just say so (and don't say "in the future, we will address weakness X").

Acknowledge funding sources. Check with your faculty advisor.

Bibliography. Job #1 is to make sure that you have been scholarly in identifying and citing related literature. But realistically, if you have less than 12 citations, it's going to look suspect (even if there were really only 12 bits of related work). You don't need more than 25 for a conference paper--- at that point, try to figure out which citations are really the most useful for the reader. No more than 15% of your citations should be citations from your own lab or close collaborators, and you should ensure that your citations include related work from the people you think are the most qualified to review your work.

Figures.
Check your figures for adequate type size, minimal reliance on color (make sure trend lines are distinguishable via both color and marker shapes, for example). Make sure axes are labeled and that axes aren't auto-scaled in a way that exaggerates results or makes results difficult to compare. Make sure aspect ratio is 1 for maps, images, etc. Note proper figure caption structure:

Figure Title (sentence fragment). Figure description-- a sentence or two (explain what the figure shows... what do the axes mean? What would a perfect result be?). Analysis-- a sentence or two (why did you show this figure? What should the reader take notice of? How does this figure support a claim?)

Figure 7. Precision/recall results on Middlebury dataset. Our method and method XYZ are evaluated on the Middlebury dataset by sweeping the detection threshold and plotting the resulting precision and recall results. A hypothetical perfect method is plotted in blue. The data shows that our method consistently outperforms method XYZ.

A cartoon sketch of typical ICRA/IROS review

Voicing. Write in complete sentences. Refer to the authors as "the authors", not by name.

Responsibility. A reviewer's job is to make sure that a paper advances the scholarly discussion and state of knowledge about a problem. This includes assessing the significance of the papers' claims and weighing the presented evidence as to whether the claims are justified. It does not mean personally attesting to the accuracy of the proofs or experimental data; a reviewer should look for problems/concerns/anomolies, but the reviewer's job is more often about *methodology* rather than specifics.

Section 1. Comprehension and Claim Evaluation. (1-2 paragraphs)

Begin with a summary of the paper. This is a reading comprehension check: did you read the paper? Did you understand what the core claims of the paper were? You don't evaluate the claims (yet)--- just make it clear that you understood them.

Is the problem domain interesting? Suppose that the claims were proven: would it be of interest to the community? (Again, you're not judging whether they've succeeded in proving their claims yet!) If the claims would be above bar, the rest of your review should focus on evaluating those claims. If the claims are below bar (e.g. maybe the problem is already solved), explain the concern and how the claims could be improved.

Section 2. Strengths and weaknesses (2+ paragraphs)

In this section, consider questions like: did the paper convincingly support their claims? Was the method principled or ad-hoc? Were there lots of parameters? Is it well written? Are the figures communicative? Could you reproduce their method? Did they design their experiments well? Did they do a good literature review? Is the problem domain interesting? Were there technical errors? In short, is it good science?

Start with paper strengths. If you can't think of something, you almost surely haven't thought hard enough.

Then discuss paper weaknesses. IMPORTANT: You should almost *always* include prescriptive advice on how the paper would be stronger. (e.g. "While the paper included a substantial evaluation on synthetic data, applying the method to a standard benchmark dataset (e.g. Middlebury) would strengthen the paper.") Don't get bogged down with minutiae ("there's a typo in sentence ....").

Section 3. Summary Assessment. (1 short paragraph).

In the aggregate, did the paper make and support interesting and useful claims? Would this paper be interesting to the community? Avoid explicitly saying "reject" or "accept"; it should be obvious how you feel based on your analysis. If the paper was hard to follow (poor English or organization), this can impact your assessment.

If the paper fell outside your area of specialty, it is appropriate to mention that here. "With the caveat that the mathematical tools used by this paper are somewhat unfamiliar to me, it seems that the proposed method is novel."

Section 4. Nits [OPTIONAL]

If there are *a few minor* corrections you'd like to offer, include them at the end of the review. Preface them with language like: "I noted a few other minor issues that could be easily addressed." Don't go nuts: you're not the author's editor. This should be stuff that could be easily fixed without needing to be re-reviewed--- e.g., "The text in figure 4 was too small for me.", or "The notation in Section IV seemed more complex than ideal; could some of the subscripts or superscripts be simplified?"

Confidential comments to the Editor

Most review processes provided a field for confidential feedback to the editor. This should be almost a copy/paste of your summary assessment, though if you want to go a bit farther on a limb saying "I'd recommend accepting/rejecting this paper", you can do that.

The Four Questions

The "four questions" are the most important questions a reviewer/audience member/proposal reviewer is likely to have. A paper answers these questions at several levels of detail: they are answered in just a couple sentences (total!) in the abstract, but might have a paragraph each in the introduction.

1. What is the problem?
Be specific, clear. Not "build a car that can drive itself", but "kino-dynamically constrained motion planning for a car"

2. Why is it important?
Why would an answer to this problem be of interest? What applications might it enable? What other fields might it impact?

3. Why is it hard?
Why is this problem worthy of a paper? Why aren't previous solutions adequate?

4. What am I doing about it?
Describe your particular approach, focusing on key technological levers. Some specificity is necessary: "applying machine learning" is not an approach. "Modeling gyro biases with a Gaussian Process" might be.