

Catchy Titles Are Good: But Avoid Being Cute

Jacob O. Wobbrock
The Information School | DUB Group
University of Washington
Seattle, WA USA 98195
wobbrock@uw.edu

ABSTRACT

The most important rule of Abstracts is that they describe the work, not the paper. Include, at most, one sentence of motivation. Save the rest of your motivation for the Introduction. Effective Abstracts focus on two things: (1) Describing what was *done*. (2) Describing what was *found* (key results). Be specific about your key findings. Instead of “many” say “84%”. Keep the Abstract to one paragraph and fewer than 200 words.

Author Keywords

Authors’ choice; of terms; separated; by semicolons; commas, within terms only; this section is required.

ACM Classification Keywords

Example: H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous;

See <http://acm.org/about/class/1998> for the full list of ACM classifiers. This section is required.

INTRODUCTION

The Introduction delivers the motivation for your paper. It explains *why* you did the work you did. This is the primary function of the Introduction.

I have found a 5-point structure for Introductions to be particularly effective. (Here, I build on advice I received as a Ph.D. student from Prof. Scott E. Hudson.)

- State of the world...
- The big BUT...
- Therefore, we did...
- The key findings are...
- The contributions of this work are...

The state of the world is a description of issues in whatever “world” is relevant to your topic. Drawing on popular press can be effective here if recent news items or data from articles support your cause.

The big BUT is where a problem is introduced. Or, similarly, a problem can be framed as an opportunity. Whether you are solving a problem or seizing an opportunity, motivate your work by connecting it to *things that matter to people*.

Copyright block. For submissions, put simply, “Submitted for review.”



Figure 1. Choose a telling figure for your paper that is placed at the top of the right-hand column on the first page. I like to make figure caption text Arial 8 pt. so that it stands apart from the body text but is small. I place my figures in-line as single “characters” in their own paragraphs, and captions as their own separate paragraphs, rather than placing either in floating text boxes, which jump around. I give captions 12 pt. after-paragraph line spacing. This figure is from [13].

One of the worst ways to motivate your work is using the “absence from the literature” argument. “Studies to date have not ...,” or “The literature is thus far silent on...,” or “Researchers have not yet examined how...” Such sentences are fine to add *after* you have established a problem or opportunity worthy of pursuit in its own right. But as stand-alone motivational statements, absence-from-the-literature does not zing. Maybe the literature is silent on an issue because that issue is not important.

After the big BUT, you will describe what you did, now in a bit more detail than in the Abstract. Devoting a paragraph to what you did is reasonable in the Introduction.

After saying what you did, you should offer the *key results* of what you did. What were the major findings? Whereas the Abstract might have devoted a sentence or two to key results, now you can devote a whole paragraph.

Finally, a good way to end your Introduction is by framing the contributions that your work makes. (See “Seven Research Contributions in HCI” [12] for examples of different research contributions.) I often structure my contributions as a numbered list within a paragraph. “The main contributions of this work are: (1) ...; (2) ...; (3) ...”

(Most papers will have one to three contributions. Work that claims to have more than three contributions is often overstating how many contributions it makes.)

By structuring your Introduction with the 5-point outline, you succinctly motivate the work against a real problem or opportunity, and hook the reader by enticing them with the key results and contributions. They will want to know more.

Often, reviewers' judgments about a piece are formed pretty quickly after reading the Introduction. If the Introduction reads poorly or is missing key aspects, that judgment can tilt negative very quickly.

Whatever expectations you set up in the Introduction, whatever research claims you make, you must deliver upon them in the rest of the paper. Over-claiming is a sure way to get your paper rejected.

Do not make the main implication of your research that more research is needed. (It can be *an* implication, but not the main one, and certainly not the one called out at the end of the Introduction, Abstract, or Conclusion.) An Introduction (or Abstract or Conclusion) that ends with, "This work opens up new directions for further research into X" as its main answer to "so what?" is weak and uninspiring. What such a statement essentially means is that the chief benefit of your research is just more research. And apparently that research is left for others to do because, for some reason, you stopped short of doing it. Instead, make the main implication of your work that the problem or opportunity you set out to address is now at least somewhat addressed. "This work paves the way for more women to be better incorporated into software development teams." "This work opens up new directions in the development of assistive technology by using adaptive user interfaces." "This work lessens the challenges faced by single mothers as they try to incorporate educational technologies into their children's lives." The point is, the outcome of your research can't just be more research, or why didn't *you* do it and gain an actual solution to report?

In a 10-page paper, it is okay if the Introduction goes a bit onto the second page. In a 4-page paper, the Introduction should probably end somewhere on the first page.

RELATED WORK

The primary function of Related Work is to answer the following four questions: "Who else has done work with relevance to this work of yours? What did they do? What did they find? And how is your work here different?" This last step is called *differentiation*.

(Sidebar: You will find that well-written Abstracts in others' papers answer "What was done?" and "What was found?" and are immensely valuable to you as you scour the literature. In contrast, Abstracts that are mostly motivation or describe their paper, not their work, require you to look into the paper itself to have any idea what was actually done. Such Abstracts are very unhelpful, but sadly, exceedingly common.)

Related Work should not read like a laundry list of who-did-what. Related Work should offer insights and education about prior work. It should *teach* readers something and help them understand the prior work better than they did before.

Skillfully written Related Work sections will often group prior work into themes. In 10-page papers, these themes may be subsections. Differentiation of the current work from prior work can then be achieved on a per-theme basis, rather than differentiating against every piece of prior work raised.

Differentiation should not be defensive. It is not required that the current work assert itself as "better" than prior work. That is for reviewers to decide. Rather, current work must be *different* than prior work. It must ask a different question, use a different method or technique, be built on different technology, or have different results. (The one exception to differentiation in this way is a pure replication study, although even then, there are usually attempts to augment or supplement the methods used.)

DESIGN OF YOUR WIDGET

After Related Work, the next sections begin the description of your work: What was *done*? If you are using design-driven research, you will often begin by describing what you created. If you have a Created Thing™, this section often takes 2-3 full pages to describe it in a 10-page paper. This section should include things like the goals for the Created Thing, the principles underlying it, the optimizations you favored, the tradeoffs you faced, the rationale for the choices you made, the design or engineering process you went through, how and why the Created Thing works, how it is built, its most important properties, its appearance, its function, how and in what situations it is used, and what its limitations are. Basically, you are trying to give a full, replicable-by-an-expert account of what you created, your Created Thing.

Figures and diagrams should be used liberally throughout this section to illustrate key points. I personally prefer figures to be inline and between paragraphs, inserted just after the paragraph that first refers to them. (The only exception to this is Figure 1.) That way, the reader encounters the figures "in the flow of reading," rather than out of sequence and having to jump their eyes back and forth from the text to the figure.

Subsections

Subsections are almost always used in these sections. Near the end of the section it is often reasonable to have an Implementation Details subsection describing things like how many lines of code are in your Created Thing (if the thing involves software), what programming language it was written in, what platform it runs on, and related. Such subsections are not particularly interesting but for technical readers, they can help to illuminate aspects of the innovation. But the main focus of these write-ups should not be on fungible implementation details; focus instead on the ideas behind the innovation and the ways those ideas were realized.

Subsubsections

Unless you absolutely must use them, avoid using subsubsections. Even the word “subsubsections” is unpleasant to read, let alone what is in them.

METHOD (BUT DON'T JUST NAME IT “METHOD”)

Regardless of whether you are reporting on design-driven research or scientific research (meaning, whether or not you have the previous section describing a Created Thing or not), you will often have a method section describing the research method(s) by which you carried out your scientific work.

(There are exceptions to this structure, such as engineering-oriented papers that provide no empirical validation of the Created Thing save the fact that they made the Created Thing *exist* in the first place. So-called “existence proof” papers are entirely legitimate; in such cases, most of those papers will be filled with descriptions of the Created Thing, and any study of that thing will be analytical, not empirical. The Method section may therefore become instead an analytical assessment of the Created Thing’s properties, rather than a description of any empirical methods.)

For papers containing empirical scientific accounts of some phenomenon, the research methods employed must be described carefully and in great detail. These research methods may be quantitative, qualitative, or mixed, and they should follow established norms in related fields such as psychology, sociology, or anthropology. The following subsections give method outlines for two common study types, quantitative laboratory experiments and qualitative field studies.

Method Outline for Quantitative Laboratory Experiments

If you are writing up your method for a quantitative laboratory experiment, a common structure found in many psychology articles is as follows:

Participants: Describe how many participants took part, their genders, their mean age ($M=32.3$ years), their age variance ($SD=6.2$ years), and any other important characteristics. If any inclusion or exclusion criteria were used, report them. The means by which participants were sampled and recruited should also be reported. If participants were compensated, say how. Give enough detail that a similar group of participants could be recruited in the future by an expert reader.

Apparatus: Describe the experiment setup, including what equipment was used, how it was arranged, how the laboratory was set up, and so on. Also describe whatever software was running and the computer system it was running on, including the operating system and version, the hardware make and model, and anything else relevant to the particular technology used, e.g., frames per second for video recording.

Procedure: Often the longest subsection within Method, the Procedure subsection should describe the process participants went through from their arrival to their departure

from the lab. What tasks did they perform? In what order? How many? What constituted a single “trial?” Enough detail should be provided that an expert reader could re-run the experiment if they were given the same participants and apparatus.

Design & Analysis: This subsection describes the experiment formally, using one paragraph to describe the experiment design in statistical terms, and one paragraph to describe the statistical analysis. An effective technique for experiments with multiple factors is to lay out the factors and their levels in a list. For example, we might write, “The experiment was a $3 \times 2 \times 2$ within-subjects design with the following factors and levels:

- *Device:* mouse, trackball, touchpad;
- *Desk:* sitting desk, standing desk;
- *Illumination:* lighted, darkness.

The analysis paragraph covers the formal statistical analysis approach used. For example, it might report: “We collected 148 data points for each of 16 participants, giving us 2368 data points in all. Twenty-two of these data points were removed due to sensor failures, resulting in 2346 viable data points for analysis. The analysis was carried out with a repeated measures ANOVA using Bonferroni adjustments for *post hoc* comparisons and the Greenhouse-Geisser correction for violations of sphericity.” As with other subsections, the test here is whether an expert reader could repeat the analysis with the description given.

Method Outline for Qualitative Field Studies

Another common type of study is the qualitative field study. Typically, such studies are reported using a wider range of structures than for laboratory experiments. Nevertheless, there are certain essential methodological aspects to cover when reporting a qualitative field study. You can also seek best practices from sociological or anthropological societies and journals.

Participants. Your method report should include coverage of the participants in the study. Depending on your study type and terminology preference, these may be informants, respondents, or just generally participants. A thick description of these participants is important, much more so than for the laboratory experiment. How did you gain access to the participants? Where and who were they? What roles did they occupy in the setting you were in? These are just a few of the questions to address about your participants.

Theory. Research built on theory may describe how the field study was designed in terms of the theoretical constructs underpinning it. Often, theory serves as an aid to decision-making when designing any kind of study, and that rationale can be invoked in the study description: Why were certain participants chosen and not others? Why was a particular field setting chosen? Why was a certain research method preferred? What does theory say about the phenomena being studied and how it should be studied?

Procedures. The procedures used in the field should also be described. These procedures may be about how the data were gathered or recorded, how informants or settings were chosen and accessed, what questions were asked, what artifacts were gathered, what probes were used, how much time was spent with certain participants or at various sites, and where and how certain observations were made. Was an ethnographic account the goal? Was participant observation or non-participant observation used? How? The one type of procedure that is usually covered separately is the procedure for data analysis, which is next.

Analysis. How were the qualitative data analyzed? How was objectivity and validity ensured? Many formal qualitative studies use grounded theory as an analysis method. If so, then the procedures that you followed for open coding, axial coding, selective coding, and theory formation should be described in detail. Give a representative sample of your codes. You could also give a link to an online version of a full coding manual. Was inter-rater reliability assessed? How? And what was the outcome? Give enough detail that the analysis could be replicated by an expert reader if he or she had your data.

RESULTS

At last we reach the exciting part, your Results! Your Results section speaks for itself: report on the results of your work in an organized way. Refrain from reporting on the *significance* of these results until the Discussion section, which comes next. For now, report on the results dispassionately. Use charts, graphs, and tables as appropriate. Good Results sections do not simply jam a bunch of numbers in, but “tell a story with data,” creating an easy-to-read narrative flow that does not make the reader do too much work to figure out what happened.

For quantitative laboratory studies, often the Results section is divided into subsections by dependent variable, i.e., the things that were measured. When including the results of statistical tests, do not just say that there was a significant effect, *interpret* that effect for the reader. Worse: “There was a significant main effect of *Illumination* on pointing time ($F(1,28)=15.79, p<.0001$).” Better: “There was a significant main effect of *Illumination* on pointing time ($F(1,28)=15.79, p<.0001$), as lighted conditions caused faster pointing than darkened conditions.”

For qualitative field studies, often the Results section is divided into subsections around emergent themes. Results sections may be quite long, incorporating numerous observations and direct quotes from participants. Organizing your results in subsections is key to creating an easy-to-follow story.

DISCUSSION

Whereas the Results section reported on what was *found*, the Discussion section reports on what was *interesting* and what *matters*. Some questions to consider when writing your Discussion section: What are the implications of your

results? What do they mean for this topic and your field? What is important and worthy of again being called to the reader’s attention? What was surprising? unexpected? intriguing? Did you fulfill the promises you set out in the Introduction via the claims you made about your work? Were any hypotheses confirmed or disconfirmed? Were the predictions of any theory upheld or refuted? What worked and what did not work? (This is a chance to be honest and show that you are objective with respect to your own work.)

The Discussion is not usually divided into subsections, but it can be if it is long. One subsection sometimes included at the end of Discussion sections is Limitations. In a single paragraph, lay out what the evident limitations of your study were and caution against faulty interpretations or over-generalization of your results.

FUTURE WORK

What are the “big idea” next steps that follow from your work? Avoid merely incremental next steps that read like a TODO list. Instead, step back and reflect on the implications of the work and where they lead. What are the *intellectual directions* that should be taken on this topic? It is often better to offer one or two well-considered inspiring next steps than a laundry list of smaller items.

CONCLUSION (NOT CONCLUSIONS)

Your Conclusion is the section in which you: (1) affirm that you have delivered on the claims made in your Introduction, (2) summarize the contributions of the work, (3) make any key points with which you would like to leave the reader, and (4) point to a bright future, a better world, for your work having been done in it.

It should not be necessary to re-report the key findings of the work, although doing so to a limited extent can be okay. The findings have already been reported, so it is better to “zoom out” and report on the key *contributions* of the work.

Worse: “We showed in this work that illuminated office conditions result in an 84% improvement in pointing speed than dark offices.”

Better: “This work contributed the first empirical study of pointing under different office illumination conditions.”

Unnecessary: “This work contributed the first empirical study of pointing under different office illumination conditions, finding that illumination improves pointing speed by 84% over dark offices.”

Try to frame the contributions of the work such that they speak to your broader scholarly community, not just those interested in your narrow topic. Imagine someone from the popular press reading your Conclusion. Could they imagine a news story on your work from what they read?

Now that you have been given numerous tips on how to structure your research paper, I am confident that you will do right by the good research you are doing by writing it in ways that others can understand and appreciate. The world will be

made better for all the work you do and the way you communicate it. Go forth and write well!

ACKNOWLEDGEMENTS

We thank all volunteers, support staff, and grandmothers everywhere for help in preparing this paper. This work was supported in part by Acme Research Corporation and the National Science Foundation under grant IIS-12345678. Any opinions, findings, conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect those of any supporter listed above.

REFERENCES

This section was written by ACM SIGCHI in their style guide for submissions to the CHI annual conference. I have kept the section as-is except for adding references [12,13].

References must be the same font size as other body text. References should be in alphabetical order by last name of first author. Example reference formatting for individual journal articles [3], articles in conference proceedings [7], books [9], theses [10], book chapters [11], an entire journal issue [6], websites [1,4], tweets [1], patents [5], and online videos [8] is given here.

Note that the `Hyperlink` style used throughout this document uses blue links; however, URLs that appear in the references section may appear in black.

1. @_CHINOSAUR. 2014. VENUE IS TOO COLD. #BINGO #CHI2016. Tweet. (1 May, 2014). Retrieved February 2, 2014 from https://twitter.com/_CHINOSAUR/status/461864317415989248
2. ACM. How to Classify Works Using ACM's Computing Classification System. 2014. Retrieved August 22, 2014 from http://www.acm.org/class/how_to_use.html
3. Ronald E. Anderson. 1992. Social impacts of computing: Codes of professional ethics. *Soc Sci Comput Rev* 10, 2: 453-469.
4. Anna Cavender, Shari Trewin, Vicki Hanson. 2014. Accessible Writing Guide. Retrieved August 22, 2014 from <http://www.sigaccess.org/welcome-to-sigaccess/resources/accessible-writing-guide/>
5. Morton L. Heilig. 1962. Sensorama Simulator, U.S. Patent 3,050,870, Filed January 10, 1961, issued August 28, 1962.
6. Jofish Kaye and Paul Dourish. 2014. Special issue on science fiction and ubiquitous computing. *Personal Ubiquitous Comput.* 18, 4 (April 2014), 765-766. <http://dx.doi.org/10.1007/s00779-014-0773-4>
7. Scott R. Klemmer, Michael Thomsen, Ethan Phelps-Goodman, Robert Lee, and James A. Landay. 2002. Where do web sites come from?: capturing and interacting with design history. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '02), 1-8. <http://doi.acm.org/10.1145/503376.503378>
8. Psy. 2012. Gangnam Style. Video. (15 July 2012.). Retrieved August 22, 2014 from <https://www.youtube.com/watch?v=9bZkp7q19f0>
9. Marilyn Schwartz. 1995. *Guidelines for Bias-Free Writing*. Indiana University Press.
10. Ivan E. Sutherland. 1963. *Sketchpad, a Man-Machine Graphical Communication System*. Ph.D Dissertation. Massachusetts Institute of Technology, Cambridge, MA.
11. Langdon Winner. 1999. Do artifacts have politics? In *The Social Shaping of Technology* (2nd. ed.), Donald MacKenzie and Judy Wajcman (eds.). Open University Press, Buckingham, UK, 28-40.
12. Jacob O. Wobbrock. 2012. Seven Research Contributions in HCI. Retrieved November 6, 2015 from <http://faculty.washington.edu/wobbrock/pubs/Wobbrock-2012.pdf>
13. Simon Kneebone and Yoland Wadsworth. 1998. What is Participatory Action Research? *Action Research International* online journal. Retrieved November 6, 2015 from <http://research-for-real.co.uk/>