How to Present Your Research Ideas

Javed Aslam *

June 25, 2011

Abstract

I have found myself consistently repeating the same advice to different students and in different contexts, and upon further reflection, I decided to codify this advice for you and for students in the future. Here it is...

1 Background

As I've mentioned to a number of you on multiple occasions, to be a successful researcher, you must (1) do good work and (2) present your work well (a) in speech and (b) in writing. Most students focus almost all of their energies on (1), but (2) is as important, if not more so. How so?

- I know many, many researchers whose work is quite good, but whose research languishes un-understood, un-implemented, and ultimately un-influential because is it was not explained well in speech or in writing. Conversely, I can count on one hand (and perhaps just one finger) the number of CS researchers who do outstanding work and are successful despite their work not being explained well in speech or writing.

- On the other hand, I know many, many researchers whose work is mediocre at best (in my opinion), but who are very successful because they explain their work well in speech and in writing.

Now the point of the above is not to say that you should do mediocre work and concentrate on explaining it well. Not at all. Rather, I want to convince you just how important it is to spend time and effort in developing your oral and writing skills.

Finally, I am not suggesting that you have to drop what you're doing and develop the verbal fluency of a Hollywood actor or the writing skills of a published novelist. Not at all. What **is** critically important is that you learn to **organize** and **present** your work in a **clear** and **coherent** manner. That is the subject of this note.

*The author is with the College of Computer and Information Science, Northeastern University, Boston, Massachusetts jaa@ccs.neu.edu

2 Venues

There are many venues and occasions in or at which one must explain their work. Here are just a few, organized roughly in increasing order of time or space allowed:

1. **The "elevator pitch":** You step into an elevator with a faculty member or colleague, and he or she asks you what you're working on; you have perhaps one minute to explain your work before the colleague gets off the elevator.

   While the elevator example is perhaps a bit fanciful, this level of research explanation happens a lot more often than you might think. For example, I often run into colleagues at conferences (or even the hallways at CCIS), and they ask me what I'm working on these days—they want a crisp, concise, and coherent explanation in just a few minutes. Another extreme example of this is the poster session at most conferences including SIGIR: You might have 50 people come by, and to each one of them, you'll have to explain your work in three minutes or less.

   It is often said that if you can't explain your work in three minutes or less to a general audience, then you really don't know what you're working on and/or you haven't thought about it deeply. Why? Well, you presumably know all the **details** of what you're working on, but you can't possibly explain those details in three minutes to a general audience. However, what you **can** explain in three minutes is the high-level description of your problem, why it's important, what high-level approach you're taking, and how your work fits into the larger context of research in the area.

   If you can't do this, crisply and succinctly, in three minutes or less, then you really need to step back and think more deeply about what you're working on and then practice...

   I can't emphasize enough how important the "elevator pitch" is—if you can do this well, then everything below is almost as simple as filling in the details.

2. **Paper abstract:** This is the written equivalent of the oral "elevator pitch"—you have a few paragraphs to explain your work.

3. **Poster board:** You have perhaps 1 meter by 1.5 meters to visually explain your work (and under three minutes to verbally explain it to everyone who comes by).

4. **Poster paper:** You have two pages to describe your work.

5. **Conference paper:** You have eight to 10 pages to describe your work.

6. **Conference talk:** You have 20 minutes to describe your work.

7. **Grant proposal:** You have 15 pages to describe your work.

8. **Journal paper:** You have 30+ pages to describe your work.
9. **PhD thesis proposal talk:** You have 45 minutes to describe your work.
10. **PhD defense / job talk:** You have one hour to describe your work.
11. **PhD dissertation:** You have 100+ pages to describe your work.

While (1) through (11) may seem radically different, I claim that how one should organize and present one's work in each of these cases is roughly the same...

3 Organizing and Presenting Research

So, how do you go about organizing and presenting research for all of the above? At the risk of oversimplification, I claim that you basically want to answer just 10 questions in five basic categories:

(1a) What problem are you working on?

(1b) Why is it important?

(2a) What have other researchers done to solve this problem?

(2b) Why is that work insufficient and/or why does the problem require further research?

(3a) What do you propose?

(3b) Why is it better and likely to work?

(4a) What results do you have?

(4b) How have you demonstrated that your work is an improvement?

(5a) What conclusions have you drawn?

(5b) What future work does this lead to?

These are the basic questions that you have to answer—it's just a matter of how much detail you give...

For an "elevator pitch" or the abstract of paper, you might allot only one or two sentences per question; for a conference paper, you might devote one section or subsection per question; for a PhD dissertation, you might dedicate a full chapter per question.

For example, here is a classic outline for a conference paper:
* Abstract (Questions 1-5)
* Introduction (Question 1)
* Related Work (Question 2)
* Methodology (Question 3)

* Results (Question 4)
* Conclusions and Future Work (Question 5)

Again, at the risk of oversimplification, it's really all about these five questions and the "elevator pitch" or abstract; the rest is details.

As a concrete example, here is an oral "elevator pitch" for my PhD dissertation:

Elevator Pitch Example

I studied machine learning and specifically classification algorithms. Consider, for example, the problem of predicting whether someone has the flu or not based upon their symptoms such as cough, body temperature, and so on: Given a reasonably large collection of training records consisting of the symptoms exhibited by flu and non-flu sufferers, one can train a machine learning algorithm to distinguish between these two groups based on their symptoms, and one can then use this trained learner to predict whether future patients have the flu or not based on their symptoms. Given the vast quantities of data collected about most every aspect of modern life, machine learning algorithms have grown more and more important, and they are used not only for medical diagnostics but also weather prediction, stock market prediction, movie and product recommendations, and a host of other important problems.

While a great many machine learning algorithms have been developed, many of them tend to be quite brittle in the sense that if even one piece of training data is mislabeled (say, "non-flu" instead of "flu"), then the machine learning algorithm will be lead down the wrong training path, producing to an inferior or failed predictor. And most unfortunately, the real world contains lots of noisy data, so this problem is critical in practice.

To address this problem, I helped developed the Statistical Query model of learning wherein a learning algorithm does not have access to individual (and possibly noisy) training records, but instead it can only ask for statistics over these training records, for example, "What fraction of flu sufferers had a high temperature?" The advantage of this model is that while detecting and correcting mislabeled individual training records is hard if not impossible—and the consequences of these mislabeled records is dire for common learning algorithms—I showed that the correct *statistics* over the training records can always be found, even in the presence of broad classes of noise. Thus, any machine learning algorithm that requires only these statistics can be made robust in the presence of noise.

I further helped to show that nearly all known machine learning algorithms can be recast into a form that requires only these statistics, and that even "weak" (poor predicting) learning algorithms can be automatically transformed into arbitrarily "strong" learning algorithms within this new model.

Thus, this work demonstrated that nearly all "weak" and "brittle" machine learning algorithms can be semi-automatically transformed into ones that are "strong" and "robust" to broad classes of noise. The challenge in practice is to

implement this model as efficiently as possible and for as many classes of noise commonly encountered.

Elevator Pitch Example Analysis

So, the question for you is... Did you understand what I did for my PhD? If not, then my elevator pitch has failed (and let me know it!). If so, then I have successfully summarized three years of my research life at MIT in 434 words and 2 minutes and 10 seconds. (I timed myself. Yes, I speak faster than most, but I bet you could read the above in under 3 minutes too.)

Note that the "elevator pitch" above has five paragraphs, roughly one each for Questions 1 through 5. I wouldn't have to modify this too much to have an abstract for a paper. I could then write that paper by expanding each paragraph to a section, and I could write a dissertation by expanding each section into a chapter.

Here's another way to think of your "elevator pitch". Imagine your work explained at all possible levels of granularity, and organize these explanations in a tree. The root level of the tree is your "elevator pitch"—it is the most concise and coherent explanation of your work. Deeper levels of the tree contain more precise and detailed descriptions of your work (and related work and future work, etc.). At the leaf-level of your tree are the most detailed and technical descriptions of your work (and related and future work, etc.).

Now imagine only describing your work at the leaf level. All the details are there, but...

You would be relying on your audience to infer what your work is all about from the details alone.

This is almost impossible for an audience to do, even if the details are explained well, and then only after some reflection. If the details are not explained well, then all is lost.

What you need to do is to **explain** what your work is all about and **only then** provide the details at whatever level of granularity is appropriate. In other words, start at the root of your tree, and work your way down. We're back to the elevator pitch...

The elevator pitch provides the framework that allows your audience to understand what you are explaining and why. It gives **them** a way to organize **your** thoughts in **their** head. You need to provide this framework; you should not rely on your audience being able to figure this out for themselves.

How is this done in practice? I'm not suggesting that you give your elevator pitch, for example, at the beginning of every talk. But, I claim, your elevator pitch should be sprinkled throughout a talk, in the form of outline slides that periodically reappear and in the form of segues that form transitions from one piece of your talk to another.

For example, if I were giving a talk on my dissertation work, and if I've had just finished explaining how and why most machine learning algorithms

are brittle to noisy data, I might flash an outline slide (which is basically just Questions 1 to 5) and then say:

"Now that we understand the problem, why it's important, and why current approaches need improvement, let me explain what I instead propose. The basic idea is that we won't allow learning algorithms to have access to the raw noisy data, but we'll instead allow them to have access to arbitrary statistics over that data. I'll show you how virtually all existing learning algorithms can be cast in such a framework and further how these statistics can be correctly inferred even in the presence of broad classes of noise. Thus, we can turn virtually all brittle learning algorithms into robust ones. Now let's see the details..."

What have I done? I've just given the audience the appropriate piece of the elevator pitch at that point in the talk while simultaneously flashing the outline slide so that they know where they are. It's something akin to a depth-first exploration of your "talk tree": At various points in your talk, you'll be finished explaining your work at some level of detail, and then you'll need to "back up" to a higher level in your talk tree, perhaps all the way to your "elevator pitch", before moving on to the next topic.

Thus, even in a one hour job talk, your three minute elevator pitch will be present: It will be scattered throughout the talk, but it is the glue that keeps the talk together and the framework that allows your audience to follow you. This is also true of research papers, concretely in the form of abstracts, but also in the form of mini-outlines and pieces of the abstract (or its ideas expressed at similar levels of detail) that reappear as segues between sections.

4 Conclusion

When it comes time to prepare a talk or write a paper (and preferably well before that), think very hard about how you could explain your work in three minutes or a handful of paragraphs. This high-level description of your work **must** be conveyed to your audience, and you cannot rely on them to infer it from the details of your talk/paper. Now organize your talk/paper along the lines suggested above. And finally, sprinkle the high-level description of your work throughout your talk/paper in the introduction, in segues between topics, and in the conclusion. Redundancy is not a bad thing when it comes to getting across your high-level ideas, and it allows your audience to follow your work. (Tell them what you are about to tell them, tell them, and tell them what you told them.)

You are much better off skipping details and getting across your high-level ideas than skipping your high-level ideas and presenting all the details. Very few members of your audience will remember the details of your talk, but you want them to remember your high-level ideas. In other words, you want them to remember the equivalent of your elevator pitch.

5 Postscript...

1. If you have suggestions on how to improve the above document, let me know.

2. The above has almost entirely focused on how to organize and present your research ideas, not on how to actually deliver them, i.e., the look of your slides, how you should speak, and the prose that you use. Patrick Winston has a very well know talk on how to deliver a talk; here is the link: `http://overstated.net/2008/01/30/patrick-winston-how-to-speak`

3. An addendum, especially for those of you who are not native English speakers: This past fall, I tried to learn Chinese while I was in Singapore, and it was brutal. I've similarly struggled with trying to learn Latin, Romanian, French, Urdu, Punjabi, and Arabic in the past. I have nothing but complete and total admiration for those of you who have picked up English as a second (or third or fourth) language. However, for better or for worse, English is the lingua franca of science, and "broken English" tends to be better tolerated in speech than in writing—reviewers can be quite picky about grammatical mistakes in papers. This is one reason why I am also quite picky about grammatical mistakes in your paper drafts; whether fair or not, it matters.