DOING RESEARCH IN THE BEHAVIORAL SCIENCES: PRACTICAL ADVICE FOR GRADUATE STUDENTS

Rudolph P. Darken
MOVES Academic Group
Department of Computer Science
Naval Postgraduate School
darken@nps.navy.mil
January 27, 2001

Introduction

For people who are schooled in technology, i.e. building things, it is difficult to break out of the box and think about technology, not from the perspective of the machine; faster, brighter, bigger, but from the perspective of the human; easier, quicker, better. However, this usually means that we have to be good at some things we aren't already good at, like how to study people. Not to worry. We can fix that. That's why we call this place a school. What follows is a very brief, and sometimes comically simplistic, presentation of how I view empirical research in the behavioral sciences. This isn't so much oriented in terms of the scientific method (whatever that is), but rather in terms of, here's a problem, how do I solve it in a meaningful way so that someone's life is made easier or safer? While most engineers will say that this is exactly how they approach their work, I like to suggest that they are the very same people who designed and built my VCR that still blinks 12:00. Enough said.

General Suggestions

- Get a research notebook.
  It is here that you will document everything you run across and every fleeting thought you have regarding your research. Even thoughts that you would consider far too immature to share with anyone should go in here. These are your private notes so no one need ever see it but you. I think I have some old grocery lists in mine. But you'll be glad you have documented everything later on because you never know when that stupid idea you had several months ago will turn out to be the breakthrough you needed.

- Read a lot.
  You need to get good at reading a lot of material, even if only abstracts, so you know who is doing what and how it all fits together. You will be the authority in your area of expertise shortly. You should know everything -- and document it in your notebook.

- Communicate.
  You are a graduate student at the Naval Postgraduate School. That makes you a part of an academic community -- even if only for a short time. That academic community is world-wide, not just here. Don't hesitate to e-mail or call people from other universities who are doing things that interest you. I get e-mail all the time from students at other universities interested in what we're doing here.

- Also, talk to each other. The best thing about being a part of a group is that we can all play a part in making each other's work better. That means taking and dishing out a lot of constructive criticism.

General Thesis Checklist

A. Select a general area of interest.
   This needs to be more specific than just "virtual environments" or "training" or whatever. This should be more like "The use of spatial sound in VEs" or "measuring spatial awareness" or something more along these lines. Note that these are still far too broad to be a thesis topic but are specific enough to move on to the next item.

B. Perform a high level literature review.
Centered on your area of interest (and be careful not to drift too far) conduct a literature review. Don't get too deep on this just yet but remain fairly high level. You can do a lot of this on the web these days but don't rely on this exclusively! You still have to go to the library for most of the substantive work. But at this early phase, the web is an essential tool for you. You are trying, at this point, to narrow your topic into a thesis topic.

How do I know when I have an area defined well enough to be a thesis topic?
Your thesis needs to get to the point that you can state it in a declarative sentence. For example: "The use of spatial sound is effective in improving performance on target acquisition tasks." In this example, you would then have to clearly state what "spatial sound" is, what "effective" means, how you define "performance", and what a "target acquisition task" is, exactly.

As you can see, if you can get to this point, you are well on your way to figuring out exactly what you have to do to prove or disprove your thesis. This is why it has to be a clear, well defined statement.

C. Finishing the literature review.
So now that you have a well defined topic, you can round out your literature review (probably tossing out a portion of what you did earlier as it is no longer relevant to your refined goals). This is where you must hit the library. I recommend starting here on campus but you may need to make a run or two up to UCSC or maybe even Stanford. Their libraries are much better than ours.

There are two reasons for this part of the literature review: (1) There is lots of material that is unavailable on the web. If you don't go to the library, you'll never find it. (2) There's a lot to be said for the reliability of your sources. Your references will probably come from the following sources. I'll list them in order of importance (reliability) most to least:

A. *Journal articles*. These are refereed and therefore the most reliable source of information and data.
B. *Books*. These are often not refereed but are excellent overviews and usually contain lots of refereed material inside them. It is not unusual for an author to compile years of work and many journal articles into a book.
C. *Conference proceedings*. These are sometimes refereed (but not always) and are shorter than journal articles. Some conferences (Like ACM CHI, SIGGRAPH, VRAIS, and others, referee the entire paper, while others like HCI International, SPIE, and others, only referee an extended abstract. Still others are by invitation only or allow marginal material in. I'm listing this because again, it says a lot about how important the information is that you might find in these sources.
D. *Magazines and trade journals*. These are typically newsstand-type periodicals rather than academic resources. Again, some sources are far better than others. An article in ACM Communications or IEEE Computer is almost like a journal article while something in Wired is not. This does not mean that these references should be skipped over. Just don't base an entire line of research on them alone.
E. *Web pages and personal communication (like a phone call)*. These are about at the same level since a person's ability to say something or put it into a web page is about the same level of effort. As you know, this material isn't filtered or refereed at all. However, there are lots of people, like myself, who put out papers that fall into the categories above, on our web pages. So it is a good idea to find out who is doing the important recent work and see if you can't find them on the web. You may find some good information.

   NOTE: The great thing about web pages is that the information can be extremely current while a journal publication is often stale by the time it hits the press. So a balance of all these sources is the best option.

At this point there may be a sub task where you know your field really well but can't determine exactly what you're going to do yet because maybe there is so little closely related work out there that you can't know what to study yet. In this case, it may be appropriate to conduct a quick exploratory study to solve this problem. You might observe people doing some task or possibly implement a formative study to collect some data to help you define your final thesis objectives. If this is necessary, you should do it. It's better to make sure you know what you're doing now than to ramble ahead without being certain. Also, clearing up any possible gray areas here will save you time later on. You won't flounder later if you make certain of what you're doing now.

4. Define exactly what you're going to do.
   If you were doing a dissertation, this would be a formal dissertation proposal document that you would have to defend publicly. For a Masters thesis, you still have to go through the same steps, but not quite so deep and you don't have to defend it to anyone but your advisor.

   This step has three essential components:

   A. Define everything that is within the scope of your thesis.
      This should now be possible at this point since you have a well defined topic and a thorough literature review completed. But you will have to determine all the smaller milestones you are going to have to hit to complete your thesis. Are you going to do an implementation? Are you going to do a study? Are you going to do both? (the most likely scenario) What will your methods be -- How will you do the implementation?, How will you do the experiment? In short, this should be a complete "todo" list of everything you have to complete, down to the smallest detail.

   B. Define everything that is not within the scope of your thesis.
      This can be considered the "cover your backside" part of the proposal phase. Not only does describing what is not your thesis help you to define what is, but it clearly states what is outside of the scope of your research so it cannot be used against you later on. You can't do everything. By stating what is and is not within the scope of your research agenda and having this agreed upon at this stage of your work, you will save yourself a lot of potential grief later on. For example, if your thesis involves training transfer of spatial sound cues in a VE, you might do an experiment in which the graphics fidelity is held constant and the aural cues are varied. You would clearly state that the relative performance on the transfer task is of interest as a dependent variable of the sound, not the video. The sound is within the scope of the research, the video is not. You might take this further to state that you are interested in only the positioning of the sound, not the timbre. So the spatial component of the sound is within the scope of the thesis while the timbre is not.

   C. Define a criteria for determining success or failure.
      This is not so much stating what you are or are not going to do, but rather in stating how we will judge your work. There absolutely must be some criteria for success. You have to state what that is. In all probability, this will be based on either the performance of the system or the performance of the operator (or team of operators), or possibly both. If you are going to build a system that can send and receive more network packets than has ever been done before, then you will state your criteria of success in terms of number of network packets over time. If you are showing that a new type of training results in better performance on some task by some group of people, then your criteria of success will be given in terms of performance measures such as time to completion, error rate, or similar.

      If you don't do this step, you leave yourself open to infinitely many possible interpretations of your work. Don't let your reader create his own success measure for you!

   At this point, you now should have a tailor-made checklist for your thesis specifically. Print it out and pin it on your wall where you can look at it regularly. Be careful not to lock yourself away and blindly try to walk through your checklist. While making progress is important, it is also important to remain in contact with the rest of the world so you can make possible alterations to your plan in

case anything changes. But try not to make any really big alterations to your plan unless you feel you have made a big mistake. And then, these changes should be carefully considered by both you and your advisor.

So this step takes us through any implementation you might have to do as well as any follow-on study. The final step in this document is the catch-all:

5. Do what you said you were going to do.
   I will only include rough directions on how to do a study here. That's a lot more complicated than I intend to get in this document. But this information should be made available to you via either a course (like OA3401, OA3402, MV4001, or MV4002) or through experience in helping actually do a study. It is best if the first study you do is not the one you will do for your thesis. Experiments are hard to do and easy to make lots of errors on. However, I think the rewards are great as it's the only way I know to make technology for people and know you've done something meaningful. Plus, working with people rather than computers exclusively can be a lot of fun.

   The basic steps for a thesis that includes an empirical study are as follows:

   A. Experimental design.
      This actually should be included in your proposal phase as it describes exactly what your experiment is. What are the treatments? What are the independent and dependent variables? What will you measure? Who will the subjects be? How many subjects will be in each group? How will you collect the data? How will you analyze the data?

      It is important to have as many people as possible look at your design before you set it in stone. It is so easy to make ridiculous mistakes simply because you got too close to the work and didn't see the big picture. You have to be able to think in terms of "What possible criticisms will the academic and military communities have regarding my experiment?" and then account for them in your design. This does not mean that you have to study every issue under the sun in your experiment, only that you have to address those most likely to be levied against you. This, again, is part of stating exactly what you are and are not going to do. For example, some people in our community are sticklers for subject pools. They might object to the use of an all male subject pool in your experiment. You might address this by stating some statistic of the male/female ratio in the military community you are working with to support your decision not to include women, or maybe use a male/female ratio in your subject pool that is statistically close to what it is for the entire population.

   B. Implementation.
      You probably have to build something to study or possibly just modify an existing system. You have to have your experimental design completed before you can complete this step (you can start it earlier, however) because you probably will have to place some data collection "hooks" in the system specifically designed for the experiment. Maybe taking times or logging positions and orientations. The basic rule of thumb here is -- Get it done. You don't get points for elegant code. I'm not suggesting that you purposely violate everything you have learned about good programming style, only that it's not what we're doing here. Good programming style is for production code. It is intended to be maintainable and have a lifecycle. You're not doing that here. Your code has a simple lifecycle. When the experiment is done, it's life is over. If you can make it work with a "goto" easier than any other way, go for it.

   C. Model building.
      As a part of some implementations, you may have to build a model of some sort or maybe get one from someplace else. We have modeling tools here to help you with this and there are people here who are experts in this area who can help you a lot. Again, get it done. However, in both this step and the actual implementation, always keep an eye on the specifications for

your experiment. Your implementation has to do what the experiment calls for -- nothing more and nothing less.

D. Pilot study.
No experiment is perfect. Yours won't be either. But the best way to make it the best it can be is to run one or more pilot subjects. You can't use this data. But going through the process of running a subject will help get the kinks out of your procedure. You have to know what to say and when. What problems might the subject have? Are your instructions clear? Have you allowed enough time for everything to happen that you need? The only way to answer these questions is to run pilot subjects and make fine adjustments to your procedure accordingly.

E. Data collection.
Now you're all set. This takes lots of time so be patient. You have to treat the last subject exactly the same as the first, even though you may be bored and tired by then. I like to videotape subjects whenever possible. This allows me to miss things while they happen knowing that I can pick them up later. I think it's important to be focused on observing your subject rather than on logging data. But sometimes this isn't an option so you have to be flexible.

F. Data analysis.
Even though you should have determined beforehand what it was you were going to analyze, almost all well designed experiments have lots of hidden items in their data that you can't possibly know about until you collect the data. So be open to this possibility. There is likely going to be some quantitative analysis consisting of plugging your data into a spreadsheet or statistical analysis tool and running calculations such as means, standard deviations, variance, t-tests, ANOVA (ANalysis Of VAriance), etc. Statistics can be used for two things: describing your data, and making inferences or interpreting your data. Both are important. Just what these test are is beyond the scope of this document but you should know something about these. This is why it's a good thing to take part in an experiment before you do your thesis research. Most importantly, don't be afraid of this. This will not be a hang-up for you.

The quantitative results answer the big questions such as "Did this training system improve performance?". But don't overlook the value of qualitative data such as behavioral observations, planning strategies, etc. These types of things can often help you to determine why something occurred, not just that it did.

G. Write-up.
The main focus of this step, of course, is the writing of your thesis. However, being a basically selfish academic, I won't be satisfied unless you produce publishable work. In fact, I'm sure you'll agree with me that if it won't be worth publishing in the end, it probably isn't worth doing in the first place. You worked hard and did something meaningful, now it's your opportunity to go to a conference or other meeting and tell people what you did. I should state clearly here that this is by no means a requirement. But I think it's important and in many ways, I think we owe it to you to let you present your work to the academic and/or military communities that motivated and guided it in the first place.

First, let me say something about the writing itself before we get into content. For many of you, this may be the first document of this type you have ever or will ever write. The biggest mistake most students make is to approach the writing of a thesis as if the fate of the world depended on it. Better to think of it as an extra-large research paper than as something bigger than it is. Take the pressure off yourself. You'll write a better document. As for style, I recommend buying a copy of Strunk and White's "Elements of Style". This is a tiny book that will probably cost you no more than $5.00 but will save you hours of time worrying about the details of technical writing. It will teach you simple rules like avoiding the first person (unlike this document) and contractions. These things are easy to do while you're writing and hard to take out after the document is written. So better to get it right the first time. Also, go to the

thesis processor and get the school guidelines to get formatting instructions. There aren't many but again, better to know up front than to wait until afterwards and have to fix it.

The thesis document looks a lot like a walkthrough of what you did. Remember, research is an incremental activity. We think about it in terms of contributions. You may have made many contributions in the course of your research. Make sure your reader is acquainted with all of them. A typical document contains sections on:

Abstract
Make sure that your abstract is more than simply baiting the reader to read the whole thing. It should concisely state what you did and what your results were. Remember, for many readers, this all they will read.

Introduction
Problem statement: Tell us what the problem is. Maybe put it into context so we can see who cares and why.
Motivation: Why is this important?, Who will benefit?
Thesis outline: Tell us what you're going to tell us.

Background
This is not a laundry list of book reports! Do not just dump your list of references and notes here. This should be organized to present the most important papers your reader has to know about with a well thought out criticism so we know why these papers are important and how they affected your work.

Approach
This is often left out or merged with the Methods section. But it's possible that you may have come up with a unique approach to what you did. Maybe you thought up a new experimental method or framework for studying your problem. This could be one of your contributions. But if you haven't done anything terribly unusual or controversial here, you can leave this section out. But if you're doing a dissertation, I would consider that highly unlikely.

Implementation
As with Approach, this may be left out if there was nothing particularly interesting with your implementation. But if you had to do something unusual in either your code development or model development, we'll want to know about it.

Methods
Tell us exactly what you did. The measure here is that the reader should be able to reconstruct everything you did from reading your document. If they can't, then you've left something out. Be very specific and tell us why decisions were made so we won't question your methods later.

Results
This is typically a presentation of the data itself. There isn't a lot of discussion as to what it means here (save that for the Discussion section). This usually looks like a lot of charts, graphs, tables, and statistical analyses that clearly present all the data you wish to present in a succinct manner.

Discussion
This is where you get to interpret the results. Tell us what it all means. What do we know from your data that we didn't know before? A very important note; Don't hide things that you can't explain or that contradict what you expected. Present it just as you would anything else. It's ok to speculate as to why something turned out like it did. Better to air your dirty laundry yourself than to have someone else find it for you. Be your own worst critic. The more you show the flaws in your own work (and it will be flawed, I can assure you), the less criticism

you are likely to take from others. They will know that you have carefully analyzed what you did and what you saw.

Conclusions
Summarize what the major findings were here in case we missed any of them. Give us a list of what all your contributions were. Tell us what you would have done differently or what improvements could be made.

Future work
Any good thesis raises more questions than it answers. This is good research. You should have uncovered several good ideas that should follow your work. This could include things that you anticipated earlier but deemed "outside the scope of your thesis" or possibly something new that you discovered during your work.

Conclusions
At this point, we throw a huge party in your honor, because you did something few people ever do -- you did research that actually meant something to somebody in a substantial way. It is my hope that this document can be helpful to you in your research endeavors, whatever you choose to do.

This document was written for the benefit of all NPS students. Any comments you might have are welcomed. Please distribute to anyone who might also benefit from reading this.

The opinions expressed in this document are my own.