How To Pass Your Qualifying Examination

(At Least In HCI and LS&T)

Mark Guzdial
With Amy Bruckman and John Stasko

January 21, 2002

As of this writing, I (Mark Guzdial) have the un-enviable notoriety of being the sole common committee member for *every* qualifying examination that has ended in a failing grade in both HCI and Learning Sciences and Technologies. One possible interpretation of that fact is that students who don't want to fail shouldn't invite me to be on their committee — but that's probably not a very useful interpretation. (Other than getting me out of some work...) A more useful interpretation is that I'm pretty well informed about what leads to an unsatisfactory outcome. In this document, I try to identify the most common mistakes and suggest how to avoid them.

1 The Point of the Qualifying Examination

The goal of the qualifying examination is to determine if the Ph.D. student is ready to do independent research and develop a proposal. What the committee is looking to see is that the student understands what research is about and how to do it. The obvious question is what we mean by "research."

The most common incorrect definition is that research is doing or making something new, unusual, and/or hard. While that description does describe research, it's not sufficient. By that definition, all start-ups are doing "research."

Research is making a contribution (in terms of new knowledge, including new interpretations or explanations of known phenomena) to the understanding of a community of researchers, then communicating that contribution. Science isn't a static thing — it's the commonly held beliefs of a community of people who are talking to one another while working on similar problems. What start-ups do *could* be research, if the people doing it could relate it to the problems of a research community and could explain it to that community in a way that helps that whole community understand the problems a bit better.

If the committee is convinced that you know this and can do this, you pass the qualifying examination. It really is that simple. Of course, the *most* convincing way to demonstrate your readiness is to already be *doing* research. The very easiest qualifying examinations for the committee is when the presentation feels more like a thesis proposal than an oral qualifying examination. But not everyone can do that, and it's certainly not necessary to be that far along to pass. It's sufficient to show us that you've done *some* research, as discussed later.

The issue of a research *community* is important. At HCI qualifiers, the committee will often discuss, "What kind of student is this? Will she most likely be talking to the UIST community? The CHI community? Maybe ICAD or UbiComp?" The reason why it's important is that each community differs in the problems they address, how they address them, and how they expect research to be explained to the community — even when the same basic "problem" or question is being addressed.

Let's take, as an example, the question of how people learn.

- Psychologists care about learning, in the context of how minds work. Psychologists typically want to know *very exactly* what's going on, so they do carefully controlled experiments in laboratories. To make sure that subjects don't have differing understandings of what's being learned, psychologists often study learning in terms of learning lists of words or non-sensical facts. There *are* psychologists who look at more complex learning outside of laboratories, but the common practice is laboratory work learning things that the subjects probably don't care about.

- Education researchers care about learning as their main focus, typically in schools whose goal is helping students become informed citizens. Education researchers work in very messy (hard to control) situations like classrooms, where some students care about the subjects and some don't. Some do laboratory work like psychologists, but the common practice is different.

- HCI researchers also care about learning, but mostly about their interfaces — sometimes in laboratories, sometimes in regular practice. Most HCI researchers aren't concerned with helping to understand how minds work, and most HCI researchers aren't concerned with what their subjects know about the *domain* of the interface (e.g., "What do bank tellers using this accounting software really know about banking?"). The common practice in HCI is to focus on learning the interface and how users do that.

Your research questions, your methods, and how you write up your results are evaluated very differently depending on the research community you're going to be talking to. If you want to talk to CHI about some cool interface you built, user evaluation is probably critical. If you want to talk to UIST about the same interface, user evaluation may not be important at all, but you better be able to talk about why it's built in an interesting way that others might be able to use — because that's what the UIST community cares about.

2 Preparing for Quals

Take your preparation for quals seriously, but remember that we're talking about learning. You can probably stay up all night and memorize lots of facts in the few days before the qual, but that won't help you with doing the synthesis needed to pass it. Take your time, work with others, think about the papers and how they relate to what you want to do. Some specific suggestions[1]:

- Pace yourself on your reading. Start as long as you can before your qual, and do a little each day. Reading for ten hours in one day will lead to your retaining little. Many people find it helpful to read first thing in the morning while their minds are fresh. Everyone has their own style–find the time of day that you are most alert. Some people find it helpful to read two technical papers *before they check their email in the morning*. After you read a paper, summarize its most important points on paper.

- Meet weekly with others studying for the qual at the same time. Bring your summaries of each paper's important points and compare them. Discuss any differences of interpretation you might have. Try to spend part of each session asking one another hard questions.

- Have a mock practice qual with other grad students in your area. Ask your friends to ask you the hardest questions they can think of, and to be as hard on you as possible. If they do their job well, when you get to the real thing it will seem easy in comparison.

[1] Bullets contributed by Amy Bruckman.

3 The Parts of the Qual, and Why They're There

The parts of the qualifying examination help the committee evaluate that basic question of whether you're ready to do research on your own and go on to define a thesis. There really is a reasonable explanation for each piece, not just that we're simply trying to run you through a variety of challenges.

3.1 The Written Qualifying Examination

The goal of the written part of the qual is to demonstrate how well you know what the community already knows. It's pretty hard for you to argue that you're adding to the existing knowledge and beliefs of a community if you don't know what the existing knowledge is! We're not really looking for encyclopedic knowledge of the field. Rather, show us that you understand what the people are talking about, especially in the areas closest to where you want to go with your thesis.

The most common mistake in the written portion is to write long lists of "Jones (1990) said... and Smith and Barney (1998) said..." It *is* important to know the key papers in the field, but we're really much more interested in knowing that you see these papers as being part of a *conversation* in a community. That's really what all research papers are: The authors' statement of what they believe is a new piece of knowledge, in response to what's been said before in this community. It's part of a conversation.

What we want to see is some synthesis of the papers in way that reflects what the community knows and believes, or perhaps captures a current dispute in the conversation. We want to see things like, "It seems that best way to do X is to use method Y (Jones, 1990), but there's also an argument to use Z when these other factors are true (Smith and Barney, 1998)." "We know that people do W and V (cite1, cite2), but it's still not clear if U is true (cite3)." Show us that you see the paper list as some of the best statements in the conversation, not as a list of facts to memorize.

You will find that committee typically is especially hard on your knowledge of the local Georgia Tech research community. Know the papers of the people here best. There are good reasons for this, besides stroking the egos of your community (and there probably is some of that...). The first reason is that we *do* talk to one another. Our papers are written as a reflection of our conversation with one another, both orally and through reading each others' papers. You're joining this local community, so you need to be an informed participant. The second reason is that we're the ones you have available to you for help as you grow as a researcher. You ought to know what we know and believe so that you know what we can help you with.

Here are John Stasko's tips on the Written exam:

- The purpose of this section is to show that you understand fundamental principles and techniques of HCI. You gain this through three things: taking the relevant courses, reading the appropriate papers, and being aware of current research topics and efforts.

- It is important that you answer a few of the questions extremely well. Exam committees are looking for students to do a very good job on at least some some of the questions. Merely having an OK, passable response on all the questions is not good enough.

- Do not get "paralysis by analysis" when reading the necessary background material. Being too locked in to the existing research, coupled with the open-book style of the exam, often leads to students answering questions by "parroting back" definitions or concepts. We already know what GOMS and heuristic evaluation are—you don't need to tell us again. What we are looking for in good answers is some synthesis of the important concepts. Critique what's relevant and what's not. Write answers that are in your own words, not just quotes from some famous researcher. If you're unsure of an answer (you see two reasonable views), explain that and give the reasons supporting each view.

- It sounds simple but... Answer the question! Carefully read what the questions is asking for and answer that. We often get long-winded responses that really don't answer what was asked. Avoid this simple, careless mistake.
- Focus on your writing skill. Even though an answer hits on important points, if it is poorly written, it will be difficult for the reviewers to extract the vital content.
- Understand what GOMS, the Model-Human Processor, and the Keystroke Model are all about. They are useful for making predictions about task performance by experts in unambiguous situations. Know when these techniques might be useful and when they definitely would not be helpful.
- Understand the differences between comparative, experimental evaluations and more qualitative, observational ones. When is each appropriate? When is each useful?
- For the User Interface software sections, understand the Model-View-Controller (MVC) concept and know all about interactor trees. (That Guzdial Guy is known for throwing an MVC question into the UIST section.)
- You should know the question(s) from your relevant research area cold, i.e., extremely well. For instance, students in the Information Interfaces group should do an excellent job answerin the Information Visualization question.
- Anticipate general evaluation and software questions that deal with some type of ubiquitous, pervasive, handheld, etc., computing (in addition to the focused UbiComp question). This area is the focus of a number of the HCI faculty here and you can expect that interest to be reflected in the exam.
- After you take the written exam, but before your orals, go over each question that you answered again and make sure that you now can give an exemplary answer to that same question.

3.2 The Portfolio

As stated earlier, the best way to show that you're ready to do research is to have done some of it. That's what the portfolio is—your demonstration that you've done *some* research. This does *not* have to be a significant, original contribution. That's what your thesis is supposed to be. What you write up for your portfolio should show that you've figured out what's known about a specific problem, and you did (a *little*) to expand our understanding of that question (a *little*).

The most common complaint that I hear about portfolios is that they're not *deep* enough. In general, that's a common complaint that researchers make about one another, so it needs some interpretation. For me, I'm looking for evidence that you really thought hard about this question:

- Did you read the literature that's related to this question? When I read the research question or problem in a portfolio, I'll typically write in the margin two or three names of people that I know have addressed this question or problem. If none of those people are mentioned, I'm worried that the author didn't go "deep" enough. If not all of them are mentioned, I typically give the author the benefit of the doubt, but I'll want to know that the author dealt with the issues (see next point).
- Did you think about how you might be *wrong*? The easiest thing to write up in research is, "I did this study. I found this fact." But it's never that simple. Somebody might argue that your study measured something different than what you thought you were measuring. Somebody else might interpret your data differently. A "deep" study considers how your study or results could be *refuted* and tries to address those points. My advisor used to tell me to be my own worst critic — how might someone argue with you, and address them explicitly in the text or

implicitly in the design of the experiment (e.g., "Someone might complain that my subjects couldn't read the instructions, but I used college students who already passed their literacy requirement.")

The most common question that I get about portfolios from people considering the qual is, "Is this enough?" The measure that the committee uses is "Is this potentially publishable, maybe in a conference somewhere?" New Ph.D. students sometimes see this as an enormous goal, but that reflects more confidence in the research community than we really deserve.

The reality is that most research barely gets read, let alone referenced, and very little of it makes a statement in the conversation that other people respond to. It's really exciting when your paper gets responded to (not just cited in passing), but it's not really all that common. That's not the measure we're applying to your portfolio. We're asking instead, "Is this interesting *enough* that someone will read it and learn something from it, where that something is relevant to this research community? Is it presented in a way that makes sense to this community?"

John's tips:

- Your portfolio should consist of your curriculum vita (academic resume), a brief (2-3 page) summary of your past research, and any papers that you have authored or co-authored.
- Remember that we're looking for research potential. Emphasize items in the portfolio that provide evidence of this.

3.3 The Oral Presentation

The oral presentation should not be a repeat of your whole portfolio. Instead we're looking for a *short* (10-15 minutes) demonstration that you really understand your problem and your solution. The *short* part is really critical here.

One of the facets of what we're looking for in the qual is knowing what's important about your story. We don't want to know, "I did this cool thing." We want to know, "In this cool thing that I did, the hard parts were... and the critical variables are..." We've hear qual presentations where the presenter has given us a bunch of measurements about their software or project, but we never knew what was important. Were all these measurements important? Which were most important and why?

For example, say that you're building a new web server. You'll probably compare it to Apache. Maybe it does better than Apache on some measures. Maybe Apache does better than your server. Does it matter? When does it matter? What makes your server interesting? Say that your server is interesting because it works entirely in RAM and can run on Palm Pilots. Then comparison to Apache's serving speed might be interesting, but the really important issues address the potential problems or implications of such a system. Does it serve fast enough to be useful? Does it serve anything interesting? What would one want to serve from a Palm, anyway?

The biggest problem that I have noted in the oral presentation is not identifying the community and the researchers in the conversation. Whose work are you responding to? Whose methods and questions are you using? Who are the people in the community that you see yourself belonging to?

Let us know about your community and whose work you're building on *right away* in your talk. This isn't like a conference talk where you can tell us your references at the end. Instead, we'd like you to contextualize what it is that you're going to tell us. A paper for CHI and a paper for UIST have very different expectations, methods, and presentation styles. We need to know how to evaluate you.

John's suggestions for the Oral Presentation:

- Plan for your research talk to be about 15 minutes, give or take a little, if you were to give it uninterrupted. You will be interrupted with questions throughout, however, and it will go on longer than that, but this is OK.

- The overall focus of your talk should be about a past/present project or projects that you have been working on. We are evaluating your prior work. This is not a proposal so you are not proposing new directions per se, but including some notion of where you are going is fine. Begin the talk with one slide on your background, where you came from, what courses you've taken, etc. Next, have a couple slides that introduce and motivate the problem you have been studying. If the committee cannot understand the problem, how can they evaluate your solutions? Next, prepare slides that discuss related, prior research. Whose work are you building upon? What has already been done in this problem area? You *must* have a solid understanding of related work. Finally, the majority of your slides should briefly describe the research that you have done and the results you have achieved. Ending with a slide or two talking about where you'd like to go in the future is good.

- Have a few slides (near the start) on Related Work. You *clearly* need to understand how your work fits into the bigger picture of all other prior, related research, and how these different pieces of work have influenced what you've done. Most importantly, don't just tick off a bullet list of related projects. That shows no synthesis. Instead, cluster related projects/work into similar themes or categories and understand how each relates to what you're doing or will be doing.

- Stay calm and be confident. Don't give wishy-washy answers. If you don't know something, say so. If you can think of a few alternatives but aren't sure which is correct, articluate them and explain your thought processes. Help the committee to see what you know and understand the relevant concepts.

3.4 Oral Follow-Up On Written Questions

The oral follow-up to your written questions is meant to give you another chance to explain more or explain better what you had written. Again, the whole game is, "Is this person ready for research?" Give us what's interesting, synthesize for us, and show us that you know your community. You may even get queries about questions that you did not answer or other questions. Be able to critique your written responses. Which answers were your best and which were your poorest? Why?

4 Conclusion

The quals may be getting tougher to pass, but not because we've decided to make them tougher. Rather, we've realized what happens if we're *not* tough at the quals time.

Your time is the most precious resource you have. If we're pretty sure that you're not going to make it to your Ph.D., it's the right thing to do to stop you. You may not realize it yourself yet. It's our job to figure out if you're not going to make it. The saddest thing in the world is to waste a few years chasing after a goal that you're not going to reach. That's what we're trying to avoid: The tragedy of wasted years.

The qual can be a *really* useful process. Students have found it valuable to get a bunch of potential-dissertation-committee members in a room and hear how they think. It's also an opportunity to hear how people think about your current research directions. Inevitably, committee members will tell you what promise they think your current research directions hold – positive or negative, that's valuable early feedback. Guaranteed, the qual will provide useful information toward your next step, the proposal.