How to get a research grant

Ian H. Witten
Computer Science, University of Calgary, Calgary T2N 1N4, Canada
Telephone (403) 220–6780; email ian@cpsc.UCalgary.CA

1. Introduction

I've just spent three years on the Natural Sciences and Engineering Research Council of Canada (NSERC) grant selection committee for Computer and Information Science. This is an arduous job, but a worthwhile—and very interesting—one. It provides an opportunity to see most of the computer science research going on in Canadian universities, and although you suffer from terrible information overload you do gain an appreciation for the breadth and excellence of the work being done. The most painful part of the job is the extraordinarily inadequate amount of money that granting agencies have to work with, and the need to reduce, or even cancel, funding for many worthwhile projects because of the extremely competitive nature of the awarding process and the dire shortage of funds.

The second most painful part of the job, which prompted me to write this article, is seeing how many capable researchers remain unfunded because they are unaware of how to write good research proposals. Many interesting projects go by the board because they are inadequately described. In the hotly competitive environment in which the grant selection committee operates, it is inevitable that inadequate or poorly-prepared research proposals receive little benefit of the doubt. The onus lies squarely on the applicant to provide clear evidence on which the committee can base a decision.

This note summarizes what I have learned about how to write research proposals, through having had to evaluate a lot of examples—good and bad—over the past three years. Provided certain mistakes are avoided, the excellence of a proposal hinges on the originality and impact of the research, and this article won't help you with that! But there are some simple guidelines that must be followed to generate a well-presented proposal.

Three principal factors are taken into account when evaluating a research grant application:

- the quality of the research program;
- the quality of the written proposal;
- the quality of the researcher.

The first factor, the quality of research, is discussed in Section 2 below. The second, the quality of the proposal, is addressed head-on in Section 3. A researcher's reputation, which is built—up or down—over time, strongly influences how his or her proposals are seen, and Section 4 gives some advice on how to present yourself in the best light. Section 5 sketches how a grant selection committee actually works. Section 6 gives some information about refereeing research grant applications, an activity that—though often seen as a chore—is absolutely essential for the health of the discipline.

This article is targetted at proposals for NSERC computer science operating grants, which are intended to provide basic support for individual researchers' work—although the same general ideas apply to any research proposal. NSERC stresses longer-term funding for individual researchers' programs more than funding for particular projects. Other granting programs have different priorities, and this should be borne in mind when preparing proposals. It should be emphasized that the views expressed here do not necessarily reflect the official policy of NSERC or any other body. Bundy (1988) has written a useful note from which some of these ideas are derived.

2. Research ideas

To do research you must formulate a *question* that your work will strive to answer. This should not just be an isolated question, but one relating to a longer-term research *theme* that evolves over a substantial part of your career—certainly much longer than the 3-year term of the average research grant. Moreover you should begin with not just a single question, but a few (although not too many) that differ in riskiness, and hence potential value. You must be able to evaluate these research questions yourself, so that you can pick good ones and present them clearly.

2.1. GENERATING RESEARCH QUESTIONS

In computer science it should not be hard to come up with good research questions. The field is young and there is much to do. Technology changes constantly, radically altering the boundaries of what is feasible, and new possibilities for research are continually opening up. There are fertile opportunities in replicating previous work more systematically and in greater depth—rational reconstruction of programs, experimental evaluation and comparison, tightening up existing conceptual frameworks, and so on. There are plenty of avenues for research in computer science!

Nevertheless, it may still be difficult to generate specific research questions. Just trying to think them up can easily lead to mental blocks. Good ideas often come from reading, discussing, explaining (and best of all, teaching) what someone else is doing. Group discussions can be fertile breeding grounds for new ideas. Read current research papers in areas that interest you, force yourself to present and explain them to others, and ideas will strike you. In my experience it's not the authors' suggestions for future research that spawn the best questions: those suggestions are ones the authors themselves haven't been able (or bothered) to pursue successfully. People who write research papers generally know far more about what they are doing than the reader, and problems that they identify but leave unsolved may well be really tough! It's better to capitalize on your more detached position to escape from the author's mind-set and think more laterally about what he's working on, rather than joining him in the tunnel of his vision and identifying open issues through his eyes.

2.2. RELATING IDEAS TO A THEME

Do not base a research proposal on just one solitary idea. Strive to give your research some breadth of scope and long-term continuity, without appearing to spread yourself too thinly. This is not easy to achieve, but merits serious effort. As months stretch into years and years into decades, your results should build up and strengthen each other so that real progress can be perceived towards answering significant and difficult questions.

An alternative research strategy is more opportunistic: to identify problems that others have formulated but failed to solve properly, and jump in with a new technique of which they are unaware and show how it can be applied. This kind of predatory strategy is often adopted by those who have special knowledge of—or an obsession with!—a particular viewpoint or tool. One danger is that to a man with a hammer, everything looks like a nail: you may be blind to the inappropriateness of your pet methodology for many of the applications you investigate. Another is that while good and plentiful results may be obtained quite quickly, over the long term the research program as a whole may take on a scrappy, uncoordinated, character. It seems better to focus your long-term efforts on particular kinds of problem than on particular kinds of solution.

2.3. SAFE VERSUS RISKY RESEARCH

Do not put all your eggs into one basket by describing only one research idea. By its very nature, it is hard to plan research, and any avenue—no matter how good it seems—may turn out to be sterile, infeasible, or simply incorrect. On the other hand, beware of promising to work on too many things, for your proposal will be criticized as being "unfocused." Reviews of proposals sometimes state explicitly that the evaluation would have been higher if fewer ideas had been included. You can spoil a good proposal by adding more to it.

Propose a mix of questions to work on; some short term and obviously answerable, others longer term, more risky, but potentially more valuable. It is important to take chances in research, and equally important to be aware of the risks being taken. Kuhn (1970) defines "normal science" as research firmly based upon one or more past scientific achievements, achievements that are acknowledged by the scientific community to supply the foundation for further practice. He contrasts this with "scientific revolutions" that question and re-structure established practice: "non-cumulative developmental episodes in which an older paradigm is replaced in whole or in part by an incompatible new one." Kuhn's distinction, which is designed for a grand scale (like Copernicus's or Einstein's revolutions in physics), also applies in miniature at the level of the individual researcher: safe versus risky research. Be aware of this distinction and propose work on different levels.

2.4. EVALUATING RESEARCH IDEAS

You have to evaluate your own ideas, assess their strengths and weaknesses, sharpen them, and present them in the most favorable light.

When you specify a goal, how will you know if you reach it? Of course, you may not expect to attain your goals, but if by chance you achieve complete success you ought to be able to tell that you have done so! Many research proposals specify goals that are so vague they could never be reached (or already have been—sometimes it's difficult to tell). It is essential to formulate goals sufficiently precisely that it will be possible to determine when they have been reached, and (if it's not completely obvious) you must explain how you will know. Goals that are stated in a way that makes it difficult to decide if they have already been achieved, or ones that are clearly completely out of reach, will destroy the credibility of any proposal.

Are your goals worthwhile, and why? The onus is on you to convince your reviewers that, if you are successful, you will have accomplished something worth doing. Of course, you might fail. But if you do succeed it is reasonable to ask what contributions will have been made to scientific knowledge (i.e. results that others can build on) or to practice (i.e. general techniques that others can use too). If you intend to prove a theorem that no-one cares about, or tackle a particular application in a way that does not shed light on others too, then research funding will be much harder to obtain (although in the latter case the application may be sufficiently useful in its own right that you can convince someone to pay for it as a development project).

Have you identified a rational approach—or, better, a few possible approaches—to tackling your chosen problem? Of course, it's very difficult to *plan* research, and that makes many people cavil at the very idea of a research proposal. But it is certainly reasonable to expect you to have some idea how to start. Obviously you should be able to say what you will do in the first few months. And you must plan something more than just "waiting for inspiration" or even "reading about the problem (and waiting for inspiration)"! Since research is evidently unpredictable and difficult to plan, have several different lines of attack in case some go wrong or do not pan out.

3. The research proposal

Given that you have the ideas, how do you describe them and make them sound worth funding? You should consider the impact of your presentation on a busy researcher—like yourself, though perhaps more experienced. You are describing your ideas to a colleague, not a business promoter. Your basic problem, as pointed out by Bundy (1988), is threefold: to convince the selection committee that

- you have identified a well-formulated goal;
- attaining this goal is a significant contribution to computer science;
- you have a good chance of reaching it with the resources requested.

3.1. DESCRIBING YOUR IDEAS

Your proposal will be evaluated by experienced, and probably sympathetic, researchers. They've been through it all themselves. They understand the difficulty of doing research and how hard it is to write a proposal. They realize that research is difficult to plan. They don't expect to be able to glean every last detail about what you want to do just by reading the proposal. But they can tell a lot about you, and the way you think, from your writing. They expect you to have thought pretty hard about your ideas, and to have worked conscientiously to explain and present them as clearly and straightforwardly as possible. They want to give you a chance, but they must justify it to themselves (and to others too). It's up to you to provide the evidence for a positive decision.

Don't make your research description a sales brochure. The kind of people who evaluate it will probably react negatively to salesmanship. On the other hand, you must make it clear that what you propose to do is worthwhile and has a good chance of success.

Acknowledge difficulties honestly. Don't try to pull the wool over the reader's eyes—he or she is probably pretty bright. If there are snags or potential problems, say so: reviewers will be impressed by your candor. If the difficulties are ones they haven't thought of, they may be impressed by your intelligence too. It is only reasonable to assume that you have thought through your proposal more thoroughly than the reviewers have; consequently if they sees problems that you don't seem to have noticed then they will be less than impressed with your efforts. It would reflect badly on your proposal if you were to describe obstacles that seem completely insurmountable, but you presumably won't be proposing work that you judge to be quite infeasible. You cannot really lose by being honest about the problems you expect to encounter.

3.2. THE RESEARCHER

As well as having good ideas, you must explain why you're fully—perhaps uniquely—qualified to carry them out. Of course, since they're *your* ideas, you automatically have a head start over others.

You must know the background to the work, the relevant literature, and what others have done. Your proposal should contain a section that reviews prior work. Space will not permit a comprehensive literature survey, and you will be unable to include many references. That makes it all the more important to select judiciously, thereby demonstrating that you have solid knowledge of the field, and the ability and good taste to make the very best use of limited space. Do not be overly introverted: mention other work besides your own. It gives a bad impression to have all (or even most) references to yourself or to a closed circle of collaborators. Avoid being involved in a small clique of researchers who publish in the same places and whose results are referred to only by one another.

For a senior researcher, the "track record" of work in the area will obviously play an important role in the evaluation of the proposal. Do not waste space by listing your own papers twice, once in the reference list and again in the personal data form or résumé. Invent a way to cross-reference from the proposal to the personal data (e.g. by numbering entries in your publication list and using letters to identify other references in the proposal). If you do not have an extensive track record, do not fret—your proposal will be judged relative to others at similar stages in their career. Everyone has to begin somewhere; the people who evaluate your proposal know that.

There is not much you can do to boost your track record, other than presenting your accomplishments fairly and accurately (see section 4). However, another most important source of information concerning whether you are the right person for the job is the understanding and insight you display when presenting and discussing the research in the proposal. The payoff for explaining your ideas clearly, eloquently, insightfully and candidly cannot be stressed too strongly.

3.3. STRUCTURE OF A PROPOSAL

Any proposal should review the context of the research, articulate the goals that will be pursued, summarize relevant prior work, describe a research plan, and give some indication of why the research is useful. Sometimes it is necessary to include a progress report on already-completed research as well.

The background should be brief and set the context for the proposal in terms of an overall research theme. The goals should project a fabric of interwoven ideas, augmenting and contributing to each other, with a mix of shorter- and longer-term, safe and risky, research, so that even if some ideas turn sour plenty will get done. One useful technique is to break down an overall goal into several interacting sub-goals or objectives—but beware of proposing too much.

For the research plan, you should at least know how you're going to start out and have some ideas for future options. Don't schedule research too firmly or too far into the future; that's unrealistic. Be prepared to describe alternative scenarios for the later stages, which hinge on how the early research turns out. Look at the problem from different points of view (theory, simulation, experimental implementations, human behavior, ...) to make it clear that you have in mind a rich variety of approaches, and the personal resources to carry them out.

Be very mindful of the need to *evaluate* your ideas, not just develop and implement them. If successful, what will be the effect of the research—how will others be able to build on the results? Will you contribute to the advancement of science, or merely develop a wonderful "look Ma, no hands" system that leaves others no better off? Sometimes such systems leave others worse off; they cannot replicate or follow up on your results and therefore cannot credibly pursue that line of research themselves.

3.4. THE PROGRESS REPORT

If you have previously been funded, you must summarize progress under earlier grants. What specific contributions have been made, where have they been published, who has taken them up, applied them, or developed them further? What students have been trained, what papers have *they* written, who has hired them? If you cannot demonstrate that you have made good use of a previous grant, your chances of getting a fresh one will clearly be diminished.

Publication delays may mean that your recent work has not yet appeared in print. Some papers based on work prior to a previously-held grant may have appeared; while this is a good sign it should not be confused with research progress stemming from the grant itself. Fortunately, your proposal will be evaluated by experienced researchers who understand the publication business and get frustrated by publication delays themselves.

3.5. PREPARING THE PROPOSAL

The people who evaluate your proposal are *busy*, even overloaded. Moreover, they are volunteers! If you can't be bothered to take the trouble to present your ideas clearly, why should they bother to read them carefully? There's a lot in it for you, much less for them. Of course, we all know that preparing research proposals is a nuisance, but reading them (by the dozen or hundred!) can be far worse. Readers will react very negatively to any signs of sloppiness in either thinking (fuzzy goals, inadequate background, unacknowledged problems ...) or presentation (poor proof-reading, spelling errors, infelicitous formatting, incomplete references ...). If you aren't sufficiently motivated or excited by your ideas to spend time honing the content and presentation of your proposal, you can't expect a sympathetic hearing from whoever is obliged to evaluate it.

Just because the reviewers are busy does not mean they will look favorably on a superficial or "popularized" proposal. Make sure there is plenty of technical content for them to pick up on. If the proposed research is highly technical, do not shy away from reflecting the technicalities in the proposal. There is nothing wrong with including a few equations if necessary, even diagrams (though be careful, especially with the latter, to ensure good use of space). Although your proposal must of necessity be brief, do not make it anaemic.

Have others read your proposal before submitting it. Encourage them to be critical, to emulate a tough reviewer, to pick out holes and ambiguities, to misunderstand where at all possible—in short, to look for ways to dislike the proposal. Probably the actual reviewers will be much more sympathetic, but you should prepare the proposal to withstand critical onslaught.

Proposals are restricted to a certain number of pages. You don't have to cover them all, but a clear exposition of complex ideas takes a certain amount of writing and most successful proposals occupy the majority of the allotted space. Don't buck the system by using a tiny typeface. Prepare the proposal in a straightforward way that won't upset the reader. It is better to get the bulk of your message across properly than to try to communicate the whole thing in detail and fail completely! Don't try to cheat by sending in more than the maximum number of pages: the proposal will be truncated before it even reaches the reviewer and the really important parts may be lost. Think of it as an exercise: part of the test is seeing how effectively you can work within specified constraints.

4. The personal data form

Along with the research proposal you will have to submit a personal data form giving information about your qualifications, the positions you have held (list them in reverse chronological order), the number of students you have supervised (specify co-supervision; divide Master's students into coursework and thesis students if applicable), your publication list, and other information. Make sure you document industrial and consulting work, along with any other "technology transfer" activity. Consider showing thesis titles and other publications by students under your supervision, listing your graduate students by name, summarizing your refereeing activity, your published reviews, and so on. What you decide to include reflects your priorities and general professionalism; it will be used by the reviewer to build a picture of you and your work.

4.1. THE PUBLICATION LIST

This is perhaps the most important part of the personal data form and you should take great care in preparing it. Gather together under separate headings papers in refereed journals (clearly indicating their "accepted" or "published" status), papers in refereed conference proceedings, other refereed items like book chapters, books, non-refereed articles, and so on. Make sure you have complete references to your papers and check that you give them the correct titles (it's surprising how many people don't!).

It is essential to be scrupulously honest when preparing the publication list. Reviewers react very negatively to any suspicion of cheating. Make sure you know for certain which of your publications are refereed. Journals, even high-profile ones, for which papers are accepted on the judgement of the editors alone are not refereed—even if one or two members of an editorial board are consulted too.

Avoid duplication in your publication list. If a conference paper was subsequently published as a book chapter, for example, choose one section in which to include it and note with that entry that it also appeared elsewhere. In general, if it's a reprint or a revision of an earlier paper, say so, and only list it once (you do gain credit from the fact that someone evidently thought it was worth reprinting!). Do not write different papers with the same (or very similar) titles.

Submitted papers should be collected together and clearly identified as such. People disagree on whether you should specify the journals to which your papers have been submitted. The argument in favor is that it gives readers a chance to judge whether you are submitting your work to appropriate places. On the other hand it might be interpreted as an attempt to glorify yourself by association with these revered organs. Moreover if your paper is not accepted you risk exposing your failure: it is surprising how much reviewers remember from one grant application to the next!

Never succumb to a temptation to mislead reviewers on the status of submitted papers—it's quite possible that someone will check with the editor of the journal and discover deception (it happens). If a paper has survived one round of refereeing and been re-submitted for a second, say so. If it has been accepted subject to minor corrections and approval by the editor, say so, giving the date of acceptance. If in doubt, spell it out.

These remarks are intended as guidelines rather than rules, and in practice there is some latitude in interpreting them. Some people prefer to list duplicate papers under each category in which they have been published, which is permissible so long as they are clearly cross-referenced. The refereed or non-refereed status of papers is sometimes not clear-cut, particularly in the case of invited papers—and ultimately, of course, it is the quality of the material that counts, not where it appears. The most important thing is to be open and honest about the status of your work. *If you are suspected of misrepresentation, your application will suffer and so will your reputation.*

4.2. ADDITIONAL MATERIAL

You may have an opportunity to submit additional material, such as preprints or reprints, to support your application. Unfortunately, reviewers are often forced to guess the quality of a paper from the journal or conference in which it appears, but if you can submit actual papers this provides a welcome opportunity to evaluate the research itself. Be sure to select reasonably recent work, and make it your best work! Do not include papers just because they have been published in prestigious journals. It may be better to choose good papers that have appeared in obscure places, or have not yet been published, as the reviewer will otherwise be quite unable to evaluate this work.

4.3. UPDATES

You may have an opportunity to submit an update to your personal data form after the application has been submitted. If you want to really annoy the reviewers, make a few minor changes to your publication list, re-sort it in a completely different order, and re-format it: then they will have to go painstakingly through the two lists to spot the differences (and determine that they are insignificant)! Although document preparation technology may make it easier for you to re-format and re-print the entire list, it will be far easier for the reader if you simply prepare a list of the differences, specifying clearly what has been changed or added, and how it should now read.

5. How a selection committee works

It helps to know a little about how a grant selection committee works. The committee has twelve or so members, each of whom reads every application in advance. Two or three members who are especially knowledgeable about the relevant research area are specifically assigned to each proposal as "internal reviewers" to evaluate it thoroughly, prepare a recommendation, and present it to the committee. The natural tendency is for internal reviewers to champion their applications where merited, and the other committee members will serve as a critical sounding-board for the presentation.

The meeting proceeds quickly. Internal reviewers say a little about your application, highlighting your credentials, what you propose to do, their evaluation, and their recommendation. If your application is any good they will be on your side, trying to persuade the other committee members of the virtues of your case. You should strive to make it easy for them! Re-read your application and imagine someone having to defend it on your behalf in the space of a few minutes. Obviously you must highlight salient points in the summary: goals, prior achievements, objectives, research plan, evaluation methodology.

Meanwhile, as your representatives present your case, the other members of the committee are leafing through the application (probably the personal data form), trying to assess the case and whether they can agree with the recommendations or not. They have studied it before, of course, but there may be hundreds of other applications and memories will need refreshing. Table 1, adapted from Bundy (1988), summarizes common reasons why proposals are rejected—bear these in mind as you prepare your proposal. There might be disagreement between the internal reviewers or with another committee member—an argument! As the discussion proceeds, the rest of the committee is silently scanning your application, listening, and thinking about it. Just imagine the impact of a poorly-prepared, scrappy proposal, and contrast it with the effect of a beautiful, tastefully-arranged document.

6. Refereeing grant applications

Selection committees depend heavily on timely and careful reviews by outside members of the research community. Each application is sent to several external referees for evaluation. Some are suggested by the applicant, others by the committee. The responses are made available to all members and referred to frequently in the committee's discussion. Indeed, in the case where no committee member has direct research expertise in the area of the

- It is not clear what question is being addressed by the proposal.
- It is not clear what the outcome of the research might be, or what would constitute success or failure.
- The question being addressed is woolly or ill-formed.
- It is not clear why the question is worth addressing.
- The proposal is just a routine application of known techniques.
- Industry ought to be doing it instead.
- There is no evidence that the proposer has new ideas that make it possible to succeed where others have failed.
- A new idea is claimed but insufficient details are given to judge whether it looks promising.
- The proposer seems unaware of related research.
- The proposed research has already been done (or appears to have been done).
- The proposer seems to be attempting too much for the funding requested and the time-scale envisaged.
- The proposal is too expensive for the probable gain.

Table 1. Some reasons for rejecting a research proposal (adapted from Bundy, 1988).

application—or when the only member who has cannot contribute because of a conflict of interest—external reviews may form the primary basis for evaluation.

Refereeing other people's applications is widely perceived as a time-consuming chore, although it can be very interesting. But just think how much more onerous it is to be on the committee itself, charged with reading hundreds of applications instead of just a handful! Ultimately it is in our discipline's interests to have the fairest possible funding decisions, and conscientious reviews play a crucial role in this. For example, NSERC evaluates the functioning of the computer and information science committee, and the computer science community at large, by the response rate to review requests: this is the kind of thing that helps whenever the committee makes requests for a larger slice of the cake. If you care about funding for computer science, you should feel obliged to contribute your share to the refereeing process.

It is important to prepare reviews thoughtfully and to the best of your ability. Unqualified praise gives the impression that you are trying to do the applicant a favor; unqualified criticism that you have a biased view. In any case it is helpful for you to summarize your previous knowledge of the applicant's work and your personal acquaintance of him or her, if any. One-line reviews give the impression that you haven't taken time to reflect upon the proposal or evaluate it properly. On the other hand, no-one wants to read a review that is longer than the proposal itself (yes, it does happen!).

The best reviewers evaluate proposals carefully and summarize the evaluation fairly, mentioning both positive and negative aspects and weighing the evidence for and against funding. Writing good reviews is just another aspect of your professionalism: it will be noticed and will enhance your reputation.

Acknowledgements

I am grateful to Rick Bunt, Brian Gaines, Saul Greenberg, Carl Hamacher and Helmut Jurgenson for making valuable comments on a draft of this article.

References

Bundy, A. (1988) "How to get an SERC grant," *AISB Quarterly 65*: 7–9.

Kuhn, T.S. (1970) *The structure of scientific revolutions.* University of Chicago Press, second edition.