How to get (and keep) a research grant*

by Ian H. Witten
revised and updated by Janice I. Glasgow

1 Introduction

Ian Witten wrote the article *"How to get a research grant"*[1]. after having just spent three years on the Natural Sciences and Engineering Research Council of Canada (NSERC) grant selection committee for Computer and Information Science. The current version of the paper was first updated by Janice Glasgow in 1996 after a similar term on the NSERC committee, and again in 1999 while she served as Group Chair for Computing and Engineering at NSERC. The updating of the document was motivated by the ever-changing climate for research funding in Canada. Because of federal budget cutbacks, getting (and keeping) an NSERC grant is becoming increasingly more difficult. Thus, the quality of a grant proposal is an even more crucial component of the funding process than it was at the time of the original article.

The job of an NSERC grant selection committee member is arduous, but worthwhile and interesting. It provides an opportunity to see most of the computer science research going on in Canadian universities, and although committee members suffer from terrible information overload they do gain an appreciation for the breadth and excellence of the work being carried out. 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

*Ian Witten is a Professor of Computer Science at the University of Waikato, New Zealand and the University of Calgary; Janice Glasgow is a Professor of Computing and Information Science at Queen's University.

[1] "How to get a research grant" initially appeared as a feature article in Canadian Artificial Intelligence, a publication of the Canadian Society for Computational Studies of Intelligence, No. 24, July 1990. The updated version was prepared in March 1996.

competitive nature of the process and the dire shortage of funds. For most Canadian researchers, an NSERC research grant (formerly referred to as operating grant) provides the core funding for their research; thus, losing this funding could significantly affect their research career.

The second most painful part of the job, which prompted Ian Witten to write the initial article, is seeing how many capable researchers remain unfunded because they are not aware of how to write good research proposals; interesting projects go by the board because they are inadequately presented. 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 the authors have learned about how to write research proposals, through having had to evaluate many such proposals – good and bad – over the years. Provided certain mistakes are avoided, the excellence of a proposal hinges on the originality and impact of the research, and this article will not help you with that! But there are some simple guidelines that should be followed to generate a well-presented proposal.

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

- the quality of the proposed research;
- the quality of the researcher;
- the training of highly qualified personnel; and
- the need for funds.

The quality of a research proposal stems from a well-planned, long-range program; this is addressed in Section 2. The quality and impact of the work must be reflected in a well-written document, as addressed in Section 3. A researcher's reputation, which is built over time, strongly influences how his or her proposal is seen, and Section 4 gives some advice on how to present yourself in the best light. It also discusses the importance of training students and researchers. 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 targeted at proposals for NSERC computer science research grants, which are intended to provide basic support for an individual researcher's work. However, many of the ideas presented apply to any research proposal. The NSERC research grants program stress long-term funding for an individual researchers' programs, rather than funding for a particular project; other granting programs may have different priorities. It should be emphasized that the views expressed here do not necessarily reflect the official policy of NSERC or any other body.

2 Research ideas

To do research you must formulate a question (or hypothesis) that your work will strive to answer (or achieve). This should not just be an isolated question, but one related to a long-term research theme that evolves over a substantial part of your career – certainly much longer than the four year term of the average research grant. Moreover, you may 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. Also, they should fit together into a coherent program with definable long- and short-term objectives.

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, i.e., 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 our experience, it is not the authors' suggestions for future research that spawn the best questions: those suggestions are ones the authors themselves have not 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 may be better to capitalize on your more detached position to escape from the author's mind-set and think more laterally about what he or she is working on.

2.2 Relating ideas to a theme

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: 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 person 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. Thus, it may be better to focus your long-term efforts on particular kinds of problems rather than on solving a string of small, weakly related puzzles.

2.3 Safe versus risky research

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 long-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 is difficult to tell). It is essential to formulate goals sufficiently precisely so that it will be possible to determine when they have been reached, and (if it is 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 apply). 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, then research funding will be difficult to obtain.

Have you identified a rational approach to tackling your chosen problem? Of course, it is tough to plan research; strategies may change depending on results you obtain along the way. But it is essential to have some idea of what methods you will apply to attempt to solve your proposed problem. You must plan something more concrete than just "waiting for inspiration" or even "reading about the problem (and waiting for inspiration)"! Since research is evidently unpredictable and difficult to plan, alternative lines of attack are sometimes useful.

2.5 Cross-discipline research

By its nature, computer science research often crosses disciplines; historically computing has often been coupled with mathematics, electrical engineering and psychology. More recently, we have seen proposals that relate computer science research with other disciplines, such as biology, mechanical engineering, linguistics, and education. In the past research that crosses disciplines has often suffered in the NSERC research grant process; grant selection committees generally favour research that promotes their own discipline, and even with evaluations from other committees it is often difficult to assess the impact and significance of work that involves multiple areas of expertise. NSERC is trying to address this issue and the computing committee is doing its best to identify and reward high-quality proposals that make contributions in more than one discipline.

It is extra tough to write a proposal that is cross-disciplinary. First, you may not know *a priori* what committee the proposal will end up in (this is determined by NSERC in consultation with the committee chairs). Thus, you may not know what audience you are addressing, e.g., a proposal written for a computing committee might have a different focus than one written for a math committee! For this reason, it is often best to focus your proposal in one discipline, and use the cross-disciplinary aspect to help demonstrate the significance and impact of the research. For example, you might propose to do research in knowledge representation, but state that the work can be motivated and evaluated in terms of cognitive psychology criteria. It is important, however, to stress the impact and significance it has on the focus discipline since it will be experts from that community who will be judging your proposal (with possibly some external advice from the secondary discipline).

NSERC does have a special committee for judging interdisciplinary proposals. However, in general your research must span at least three disciplines to be eligible for consideration by this committee. This document does not address such research.

Despite the problems discussed above, the computing committee does recognize the importance of research that crosses borders. It certainly strengthens your own research if it can be demonstrated that it has significance in other areas. Just be careful that you are not just applying known technology to a new problem domain: your proposal will be judged primarily on how you are contributing to the advancement of computer science.

3 The research proposal

Carefully read and follow all instructions provided with the grant application. Your application may be disqualified if you do not follow specifications, such as font size and format. As well, committee members do not appreciate reading material presented in a non-standard way – and the last thing you want to do is make an overextended reviewer unhappy!

Given that you have the ideas, how do you describe them and make them sound worth funding? First, remember that you are describing your ideas to a colleague, not a business promoter: be positive and optimistic about your work, but avoid making a sales pitch. Your basic problem, as pointed out by Bundy (1988), is threefold. It is to convince the selection committee that:

- you have identified a well-formulated goal;
- attaining this goal would be a significant contribution to computer science;
- you have a good chance of attaining the goal with the resources available.

One of the complexities of writing a research proposal is that you have to address two audiences: 1) the internal and external reviewers, who are likely to be experts in your field of interest, and 2) the remainder of the committee, who are computer scientists but may have limited knowledge of the area in which you are working. Your proposal must have something for both audiences; there should be enough depth and detail to please the expert, but you must also convince the non-expert of the importance and impact of your proposed research. In particular,*the abstract should be written for the general computer science audience.*

3.1 Describing your ideas

Your proposal will be evaluated by experienced, and probably sympathetic, researchers. They have been through it themselves and understand the difficulty of proposal writing and conducting research. They realize that research is difficult to plan and do not 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). It is up to you to provide the evidence for a positive decision.

Do not 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: do not try to pull the wool over the readers' eyes. If there are snags or potential problems, say so; reviewers will be impressed by your candor. If the difficulties are ones that they have not 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 see problems that you do not 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 will not be proposing work that you judge to be 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 are fully – and perhaps uniquely – qualified to carry them out. Of course, since they are your ideas, you automatically have a head start over others.

You must know the background for the work, i.e., the relevant literature in the field and how it relates to your research. Your proposal should contain a brief ($\leq$ 1 page) section that reviews this prior work.[2] 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. Make sure your literature review is up-to-date, including recent publications in the area if they exist, and if not, consider explaining why (lest your proposal be seen as belonging to a bygone era).

For a senior researcher, the "track record" of work (especially recent) 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 (PDF) or resume. Invent a way to cross-reference from the proposal to the PDF (e.g. by numbering entries in your publication list and using letters to identify other references in the proposal). If you are a relatively junior researcher who does not have an extensive track record, do not fret – your proposal will be judged relative to others at similar stages in their career. New applicants are evaluated primarily on their potential. In these cases, where there is generally little track record, the quality of the proposal is even more crucial. Those who are being considered for renewal for the first time are judged on whether they have begun to develop an independent research program. Publications that are based only on research done during a PhD or postdoc may not convince the committee of this.

When considering subsequent renewals, evidence of an independent long-term research program is essential. This does not mean that team research is not deemed important, quite the contrary: recent communications from NSERC have emphasized that collaborative and concerted activities should be actively encouraged and that grant selection committees should give credit

[2] Some applicants prefer to combine the literature review with the description of proposed work. This is allowable and sometimes more effective than putting it in a separate section.

to effective research interactions. This was not always perceived to be the case in the past, when it was considered that more emphasis was placed on sole authored papers and there was sometimes a lack of appreciation for strong multi-authored papers. The current sentiment is that creativity and innovation are at the heart of all research advances, whether made individually or as part of a group effort. In an NSERC research grant proposal it is useful for the applicant to clearly define his/her personal contributions to joint research (there is a place for this in the PDF).

There is not much you can do to boost your track record, other than presenting your accomplishments fairly and accurately (see Section 4). However, an 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. A progress report on completed research is also required for renewals.

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 short- and long-term, safe and risky, research. One useful technique is to break down an overall goal into several interacting sub-goals or objectives – but beware of proposing too much. Also ensure that sub-goals fit together into a cohesive research program, rather than what might appear to be a collection of unrelated projects.

The majority of the proposal should be devoted to a careful description of the research objectives and the methodology whereby these objectives will be achieved. For the research plan, you should at least know how you are going to start out and have some ideas for future options. Do not schedule research too firmly or too far into the future; that is unrealistic. Be prepared to describe alternative scenarios for the later stages, which hinge on how the early research turns out. It might be useful to look at the problem from different points of view (theory, simulation, experimental implementations, human behavior, etc.) provided you have the background and resources to carry this out. Be 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 they 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 if they cannot replicate or follow up on your results.

3.4 The progress report

If you have previously been funded, you must summarize the progress made under the previous grant: What specific contributions have been made, where have they been published, and who has taken them up, applied them, or developed them further? If you cannot demonstrate that you have made good use of a previous grant, the chances of your grant getting renewed will clearly be diminished.

Publication delays may mean that your recent work has not yet appeared in print. Some papers based on work related to research funded from other sources may also have appeared; while this is a good sign, it should not be confused with research progress stemming from the NSERC grant itself. Fortunately, your proposal will be evaluated by experienced researchers, who understand the publication business (particularly the fact that ambitious research projects may require multiple sources of funding) and that delays in publication often occur.

Be careful of your balance between progress and proposed research. In the past, a senior researcher might have succeeded with a proposal that mostly describes past work and then says something to the effect of "I plan to continue doing more or less the same thing". In the current competitive environment, this no longer works. Even those with an excellent track record must include a detailed and well-written research proposal for the work they plan to do.

Note that some of the items that indicate progress have a natural place in the PDF. Use the space allotted for describing your proposal wisely.

3.5 Budget

One of the criteria for a research grant is need. The unfortunate circumstance is that almost all researchers "need" more money than the committee can possibly provide. Despite this, budgets should be carefully prepared and justified. Most importantly, be realistic: some applicants have lost credibility by proposing unrealistic budget items. For example, if you are a new researcher then it might be difficult (but not impossible) to justify funding for a postdoc. Similarly, if you are a senior researcher who has not supervised graduate students recently then a budget that includes support of several students might be questioned.

Make certain that your budget items are allowable by NSERC. Since the guidelines change year to year, check thes for each new application.

A budget should reflect the actual cost of carrying out your proposed research – even if such funding is not available from the granting agency. Note that the committee cannot give you more money than you ask for, so do not be too modest in what you request!

Researchers are often funded from multiple sources, suggesting that there may be some question of need. In general, additional sources of funding are seen as a positive reflection on the researcher and his/her work. However, there are cases where the amount of time that the applicant has for the new proposal might be a consideration. It is important that you make clear how work being proposed is unique and distinct from research funded from other sources.[3]

One major budget item affecting the need for funds is the support of graduate students. For those who are not supporting students, either by choice or by circumstance (e.g., there is no graduate program at their institution), the perceived need for funds is typically less than for those supporting students.

3.6 Preparing the proposal

It is important to take great care to present your ideas clearly: the people who evaluate your proposal are busy, even overloaded, volunteers. Thus, they will probably react negatively to any signs of sloppiness in either thinking (fuzzy goals, inadequate background, unacknowledged problems, etc.) or presentation (poor proof-reading, spelling errors, infelicitous formatting, incomplete references, etc.). If you are not sufficiently motivated or excited by your ideas to spend time honing the content and presentation of your proposal, you cannot expect a sympathetic hearing from whoever is obliged

[3]This should be done on the supplementary budget page, rather than in the research proposal itself.

to evaluate it.

Reviewers do not generally look favorably on superficial or "popularized" proposals. 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).

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 more sympathetic, but you should prepare the proposal to withstand a critical onslaught. Also, if possible, read other proposals – particularly those that have been successful in the past – to pick up clues of how your presentation could be improved. Some departments and faculties have research committees or officers that provide assistance in preparing proposals; take advantage of any such resources that are available to you. It will probably be four years before you have another chance to apply for a grant, so it takes a long time to recover from funding cuts resulting from a poorly-prepared proposal. This is particularly true for first-time applicants, where the implication of a good versus a bad proposal is funding versus non-funding.

Proposals are restricted to a certain number of pages. You do not 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. Do not buck the system by using a tiny typeface (chances are your proposal will become ineligible for funding if it does not conform to the proposed standards). Prepare the proposal in a straightforward way that will not 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! Do not try to cheat by sending in more that the maximum number of pages as the proposal will be truncated before it even reaches the reviewers and 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 (PDF)

Along with the research proposal, you will have to submit a PDF giving information about your qualifications, the positions you have held (list them in reverse chronological order), the number of students you have supervised, your publication list, and other information. Make sure you document industrial and consulting work, along with any "technology transfer" activity. Consider showing thesis titles and publications by students under your supervision, listing your undergraduate and graduate students and post-docs by name (along with their career progress), 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 PDF 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 provide correct and complete references to your papers.

It is essential to be scrupulously honest when preparing the publication list. Reviewers react negatively to any suspicion of cheating. Make sure you know for certain which of your publications are refereed. Only list publications for the specified time period.

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 is 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!). Avoid writing 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.

Never succumb to the temptation to mislead reviewers on the status of submitted papers – it is 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 both conference papers and journal papers that are straightforward extensions of them separately, 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.

Justify the journals and conferences where you are publishing, particularly if they are not ones that are easily recognizable by members of the committee. For example, interdisciplinary work might be best presented in forums that are accessible by both disciplines.

4.2 Most significant research contributions

There is a fair bit of flexibility in what you list as your most significant contributions. These may be the papers that you feel have had the most impact. It also may include software that has been developed, key ideas (that may be supported by multiple papers) or books that have arisen from your research ideas. You should make it clear why these contributions were significant. This can be achieved by noting others who have followed up on the work (using your research as a basis), applications that have arisen from the research, invitations to present the results in seminars, etc.

4.3 Explanation of research contributions

Make clear what your personal contribution is to joint publications. If a publication is primarily the work of a graduate student under your supervision – say so. If the order of authors reflects (or does not reflect) their respective contribution – say so. If anything is not clear, then a reviewer might assume the worst case scenario.

4.4 Contributions to the training of highly qualified personnel

The training of students and researchers is an important criterion in the assessment of a research grant proposal. As with publications, the focus is not just on quantity, but also on quality: one graduate student who goes on to be a first class researcher is generally preferable to three unemployed graduates.

Do not just list numbers: list thesis titles, say where students went after graduating, mention awards or publications they may have received under your supervision, list co-supervisors, etc. Doing so indicates that the training of students and researchers is a high priority to you.

The merit of training is judged in conjunction with the quality of the researcher. Someone who has a poor track record, or a badly presented proposal, may not be considered the best person to train new researchers. In such cases having supervised a large number of students might be thought of as a negative, rather than a positive, contribution.

The research training of undergraduate students (particularly at universities that do not have a graduate program), is considered as a worthwhile contribution. This is especially the case if the students are encouraged to go on to graduate school after this experience.

4.5 Other evidence of impact

This section of the PDF allows you to present further evidence to the significance of your research. In particular, it is useful to list awards and honours related to your work. In some cases you might wish to describe the significance of the honour if it is not a well known one, e.g., a listing of a best paper award could include the number of papers submitted and accepted to the conference.

Prestigious invited talks also support the significance of your research. Positions on program committees and journal boards suggest that your peers value your contributions.

One obvious evidence of impact in computer science is the dissemination of software resulting from research. If you have written programs that are being used by other researchers or industry, this could be an important contribution. Consulting activities are also worth mentioning, if they relate to your research.

4.6 Delays in research

If there have been any significant reasons for delay in research, then it is may be useful to specify these in the PDF. The committee is sympathetic to the fact that that there may be situations that may result in decreased research productivity, and this can be taken into consideration when evaluating your research record. Situations such as maternity leave or illness are valid reasons for a temporary slow down in activity. Administrative responsibilities may also have impaired your progress. It is important, however, that these delays be perceived as temporary (e.g., if your research has slowed down because you have had a term as an associate dean and have subsequently been appointed to a term as dean, then there is no reason to think that the situation will improve). Also, be careful not to cite delays that are perceived as a normal part of a researcher's responsibility (e.g., graduate student advising or program committee membership).

4.7 Additional material

You may have the 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 contributions more closely. 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.

5 How a grant selection committee works

It helps to know a little about how a grant selection committee works. Following is a description of how the 1998/99) Computing and Information Science grant selection committee operated; other NSERC committees tend to have similar procedures, but we cannot guarantee this.

Previously, computing has only one committee that handles all applications (e.g., Math is split into two subcommittees).[4] The proposals reviewed by the computing grant selection committee come mainly from computer science departments of Canadian universities; applications may also come from other departments, such as business, library science, math and electrical engineering. Certain research institutes (e.g., CRIM) and colleges have employees who are eligible for funding. As well, researchers in industry who are adjunct professors at a university may be eligible to apply for a grant.[5]

Each application is read carefully by members of the committee (this may be the entire or some large subset of the committee). Two or three members, who are especially knowledgeable about the relevant research area, are specifically assigned to each proposal as "internal reviewers" to evaluate the proposal thoroughly, read the additional material, and prepare and present a recommendation to the committee. The natural tendency is for internal reviewers to champion their applications where merited, and the other committee members serve as a critical sounding-board for the representation. The meeting proceeds quickly – on average, there is less than ten minutes available per application (though in cases where there is disagreement, the discussion could last longer). Internal reviewers summarize the application, highlighting the applicants credentials, what they propose to do, their evaluation, and finally their recommendation. If they judge the proposal to be good then they will act as a proponent, trying to persuade the other committee members of the virtues of the 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.

Proposals are considered in categories: all new applicants are handled first, then applications from researchers whose current grants fall into the same funding range (e.g. \$0 - \$20,000) are considered consecutively.

The recommendation of an internal reviewer includes a proposed amount

---

[4] There is currently a proposal to break the computing committee into two subcommittees beginning in 1999/00. If this is the case then the subcommittees should individually operate in much the same way as the current larger committee.

[5] Eligibility requirements vary from year to year, so check with NSERC if there is any uncertainty about the current policies.

for funding. Each committee member must balance a budget, thus a recommendation to increase the funding of one applicant might imply the decrease or non-funding of another grant application.

As your representatives present your case, the other members of the committee are reviewing the application and external referee reports trying to assess the case and decide 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 – maybe even an argument (it has happened)! 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.

The final recommendation for funding is done by a "Dutch auction" method. An initial funding amount is proposed; this is generally the highest of the amounts suggested by the internal reviewers, but a higher amount can be recommended by another member of the committee if they feel it is warranted. Voting then proceeds where only those who are designated readers (recall that there are nine readers and five non-readers for each proposal) have the right to vote. The amount incrementally decreases (usually by $1,000 decrements) until a majority of eligible voters have their hands up. Anyone who is in conflict with a proposal (e.g., at the same university or a co-author) is a non-reader and must leave the room during the discussion and voting process. The final amount decided upon is only a recommendation to NSERC; given that the overall budget may be altered after competition week, proposed amounts may vary from those eventually awarded.

6 Refereeing grant applications

Selection committees depend heavily on timely and carefully prepared 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 deliberations.

Refereeing other people's applications is widely perceived as a time-consuming chore, although it can be interesting. 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 (such as in the reallocation exercise that occurs every four years). If you care about the quality of computer science research and future funding for the discipline, 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 suggests 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 have not 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 that the proposal itself (yes, it does happen!). The best reviewers evaluate proposals carefully and summarize the contributions 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.

7 Conclusion

No amount of care and effort in preparing a research grant proposal will compensate for a weak research program. However, a poorly prepared proposal can prevent a strong research proposal from being funded at the level it deserves. The authors of this document hope that a better understanding of the NSERC review process can assist researchers in having a clearer idea of how to best present their ideas and contributions for future research grant competitions.

- It is not clear what question or hypothesis 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 methodological details are given to judge whether it looks feasible.
- 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)

Acknowledgements

Ian Witten is grateful to Rick Bunt, Brian Gaines, Saul Greenberg, Carl Hamacher and Helmut Jurgensen for making valuable comments on a draft of the original article. Janice Glasgow thanks Uri Ascher for his insightful comments on a draft of the revised version of the document, and Nancy Barker for her assistance in recreating the initial document.

References

Bundy, A. (1988) "How to get a SERC grant," AISB Quarterly 65: 7-9.

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