Chapter 2
Getting Started

Science is more than a body of knowledge; it is a way of thinking.
Carl Sagan
The Demon-Haunted World

There are as many scientific methods as there are individual scientists.
Percy W. Bridgman
On "Scientific Method"

There are many ways in which a research project can begin. It may be that a conversation with a colleague suggested interesting questions to pursue, or that your general interest in a topic was crystallized into a specific investigation by something learnt in a seminar, or that enrollment in a research degree forced you to identify a problem to work on. Then definite aims are stated; theories are developed or experiments are undertaken; and the outcomes are written up.

The topic of this chapter is about getting started: finding a question, working with an advisor, and planning the research. The perspective taken is a practical one, as a working scientist: What kinds of stages and events does a researcher have to manage in order to produce an interesting, valid piece of work? This chapter, and Chaps. 3, 4, 14, and 15, complement other parts of the book—which are largely on the topic of how research should be described—by considering how the content of paper is arrived at.

Thus this chapter concerns the first of the steps involved in doing a research project, which broadly are:

- Formation of a precise *question*, the answer to which will satisfy the aim of the research.
- Development of a detailed *understanding*, through reading and critical analysis of scientific literature and other resources.
- Gathering of *evidence* that relates to the question, through experiment, analysis, or theory. These are intended to support—or disprove—the hypothesis underlying the question.
- Linking of the question and evidence with an *argument*, that is, a chain of reasoning.
- Description of the work in a publication.

J. Zobel, *Writing for Computer Science*, DOI 10.1007/978-1-4471-6639-9_2

Learning to do research involves acquisition of a range of separate skills. It takes experience to see these skills as part of a single integrated "process of research". That is, many people learn to be researchers by working step by step under supervision; only after having been through the research process once or twice does the bigger picture become evident.

Some newcomers try to pursue research as if it were some other kind of activity. For example, in computer science many research students see experimentation as a form of software development, and undertake a research write-up as if they were assembling an essay, a user manual, or software documentation. Part of learning to be a scientist is recognition of how the aims of research differ from those of coursework.

A perspective on research is that it is the process that leads to papers and theses, because these represent our store of accepted scientific knowledge. Another perspective is that it is about having impact; by creating new knowledge, successful research changes the practices and understandings of other scientists. Our work must be adopted in some way by others if it is to be of value. Thus another part of learning to be a scientist is coming to understand that publication is not an end in itself, but is part of an ongoing collaborative enterprise.

Beginnings

The origin of a research investigation is typically a moment of insight. A student attending a lecture wonders why search engines do not provide better spelling correction. A researcher investigating external sorting is at a seminar on file compression, and ponders whether one could be of benefit to the other. An advisor is frustrated by network delays and questions whether the routing algorithm is working effectively. A student asks a professor about the possibility of research on evaluation of code reliability; the professor, who hadn't previously contemplated such work, realises that it could build on recent advances in type theory. Tea-room arguments are a rich source of seed ideas. One person is idly speculating, just to make conversation; another pursues the speculation and a research topic is created. Or someone claims that a researcher's idea is unworkable, and a listener starts to turn over the arguments. What makes it unworkable? How might those issues be addressed?

This first step is a subjective one: to choose to explore ideas that seem likely to succeed, or are intriguing, or have the potential to lead to something new, or that contradict received wisdom. At the beginning, it isn't possible to know whether the work is novel or will lead to valuable results; otherwise there would be no scope for research. The final outcome is an objective scientific report, but curiosity and guesswork are what establish research directions.

It is typically at this stage that a student becomes involved in the research. Some students have a clear idea of what they want to pursue—whether it is feasible, rational, or has research potential is another matter—but the majority are in effect shopping for a topic and advisor. They have a desire to work on research and to be creative, perhaps without any definite idea of what research is. They are drawn by a particular area or problem, or want to work with a particular individual. Students may talk through a range of possible projects with several alternative advisors before making a definite choice and starting to work on a research problem in earnest.

Shaping a Research Project

How a potential research topic is shaped into a defined project depends on context. Experienced scientists aiming to write a paper on a subject of mutual interest tend to be fairly focused: they quickly design a series of experiments or theoretical goals, investigate the relevant literature, and set deadlines.

For students, doing a research project additionally involves training, which affects how the work proceeds. Also, for a larger research program such as a Ph.D., there are both short-term and long-term goals: short-term goals include the current specific explorations, which may be intended to lead to an initial research paper; the long-term goals are the wider investigation that will eventually form the basis of the student's thesis.

At the beginning of a research program, then, you need to establish answers to two key questions. First, what is the broad problem to be investigated? Second, what are the specific initial activities to undertake and outcomes to pursue? Having clear short-term research goals gives shape to a research program. It also gives the student training in the elements of research: planning, reading, programming, testing, analysis, critical thinking, writing, and presentation.

For example, in research in the 1990s into algorithms for information retrieval, we observed that the time to retrieve documents from a repository could be reduced if they were first compressed; the cost of decompression after retrieval was outweighed by savings in transfer times. A broad research problem suggested by this topic is whether compression can be of benefit within a database even if the data is stored uncompressed. Pursuing this problem with a research student led to a specific initial research goal: given a large database table that is compressed as it is read into memory, is it possible to sort it more rapidly than if it were not compressed at all? What kinds of compression algorithm are suitable? Success in these specific explorations leads to questions such as, where else in a database system can compression be used? Failure leads to questions such as, under what conditions might compression be useful?

When developing a topic into a research question, it is helpful to explore what makes the topic interesting. Productive research is often driven by a strong motivating example, which also helps focus the activity towards useful goals. It can be easy to explore problems that are entirely hypothetical, but difficult to evaluate the effectiveness of any solutions. Sometimes it is necessary to make a conscious decision to explore questions where work can be done, rather than where we would like to work; just as medical studies may involve molecular simulations rather than real patients, robotics may involve the artifice of soccer-playing rather than the reality of planetary exploration.

In choosing a topic and advisor, many students focus on the question of "is this the most interesting topic on offer?", often to the exclusion of other questions that are equally important. One such question is "is this advisor right for me?" Students and advisors form close working relationships that, in the case of a Ph.D., must endure for several years. The student is typically responsible for most of the effort, but the intellectual input is shared, and the relationship can grow over time to be a partnership of equals. However, most relationships have moments of tension, unhappiness, or disagreement. Choosing the right person—considering the advisor as an individual, not just as a respected researcher—is as important as choosing the right topic. A charismatic or famous advisor isn't necessarily likeable or easy to work with.

The fact that a topic is in a fashionable area should be at most a minor consideration; the fashion may well have passed before the student has graduated. Some trends are profound shifts that have ongoing effects, such as the opportunities created by the Web for new technologies; others are gone almost before they arrive. While it isn't necessarily obvious which category a new trend belongs in, a topic should not be investigated unless you are confident that it will continue to be relevant.

Another important question is, is this project at the right kind of technical level? Some brilliant students are neither fast programmers nor systems experts, while others do not have strong mathematical ability. It is not wise to select a project for which you do not have the skills or that doesn't make use of your strengths.

A single research area can offer many different kinds of topic. Consider the following examples of strengths and topics in the area of Web search:

Statistical. Identify properties of Web pages that are useful in determining whether they are good answers to queries.

Mathematical. Prove that the efficiency of index construction has reached a lower bound in terms of asymptotic cost.

Analytical. Quantify bottlenecks in query processing, and relate them to properties of computers and networks.

Algorithmic. Develop and demonstrate the benefit of a new index structure.

Representational. Propose and evaluate a formal language for capturing properties of image, video, or audio to be used in search.

Behavioural. Quantify the effect on searchers of varying the interface.

Social. Link changes in search technology to changes in queries and user demographics.

As this list illustrates, many skills and backgrounds can be applied to a single problem domain.

An alternative perspective on the issue of how to choose a topic is this: most projects that are intellectually challenging are interesting to undertake; agonizing over whether a particular option is *the* project may not be productive. However, it is also true that some researchers only enjoy their work if they can identify a broader value: for example, they can see likely practical outcomes. Highly speculative projects leave some people dissatisfied, while others are excited by the possibility of a leap into the future.

When evaluating a problem, a factor to consider is the barrier to entry, that is, the knowledge, infrastructure, or resources required to do work in a particular field. Sometimes it just isn't possible to pursue a certain direction, because of the costs, or because no-one in your institution has the necessary expertise. Another variant of the same issue is the need for a codebase, or experience in a codebase; if investigation of a certain query optimization problem means that you need to understand and modify the source code for a full-strength distributed database system, then possibly the project is beyond your reach.

As research fields mature, there is a tendency for the barrier to entry to rise: the volume of background knowledge a new researcher must master is increased, the scope for interesting questions is narrowed, the straightforward or obvious lines of investigation have been explored, and the standard of the baselines is high. If a field is popular or well-developed, it may make more sense to explore other directions.

Project scale is a related issue. Some students are wildly ambitious, entering research with the hope of achieving something of dramatic significance. However, major breakthroughs are by definition rare—otherwise, they wouldn't be major—while, as most researchers discover, even a minor advance can be profoundly rewarding. Moreover, an ambitious project creates a high potential for failure, especially in a shorter-term project such as a minor thesis. There is a piece of folklore that says that most scientists do their best work in their Ph.D. This is a myth, and is certainly not a good reason for tackling a problem that is too large to resolve.

Most research is to some extent incremental: it improves, repairs, extends, varies, or replaces work done by others. The issue is the magnitude of the increment. A trivial step that does no more than explore an obvious solution to a simple problem—a change, say, to the fields in a network packet to save a couple of bits—is unlikely to be worth investigating. There needs to be challenge and the possibility of unexpected discovery for research to be interesting.

For a novice researcher, it makes sense to identify outcomes that can clearly be achieved; this is research training, after all, not research olympics. A principle is to pursue the smallest question that is interesting. If these outcomes are reached early on, it should be straightforward, in a well-designed project, to move on to more challenging goals.

Some research is concerned with problems that appear to be solved in commercial or production software. Often, however, research on such problems can be justified. In a typical commercial implementation the task is to find a workable solution, while in research the quality of that solution must be measured, and thus work on the same problem that produces similar solutions can nonetheless have different outcomes. Moreover, while it is in a company's interests to claim that a problem is solved by their technology, such claims are not easily verified. In some cases, investigation of a problem for which there is already a commercial solution can be of as much value as investigation of a problem of purely academic interest.

Research Planning

Students commencing their first research project are accustomed to the patterns of undergraduate study: attending lectures, completing assignments, revising for exams. Activity is determined by a succession of deadlines that impose a great deal of structure.

In contrast, a typical research project has just one deadline: completion. Administrative requirements may impose some additional milestones, such as submission of a project outline or a progress report, but many students (and advisors) do not take these milestones seriously. However, having a series of deadlines is critical to the success of a project. The question then is, what should these deadlines be and how should they be determined?

Some people appear to plan their projects directly in terms of the aspects of the problem that attracted them in the first place. For example, they download some code or implement something, then experiment, then write up. A common failing of this approach to research is that each stage can take longer than anticipated, the time for write-up is compressed, and the final report is poor. Yet the write-up is the only part of the work that survives or is assessed. Arguably, an even more significant failing is that the scientific validity of the outcomes can be compromised. It is a mistake, for example, to implement a complete system rather than ask what code is needed to explore the research questions.

A strong approach to the task of defining a project and setting milestones is to explicitly consider what is needed at the end, then reason backwards. The final thing required is the write-up in the form of a thesis, paper, or report; so you need to plan in terms of the steps necessary to produce the write-up. As an example, consider research that is expected to have a substantial experimental component; the write-up is likely to involve a background review, explanations of previous and new algorithms, descriptions of experiments, and analysis of outcomes. Completion of each of these elements is a milestone.

Continuing to reason backwards, the next step is to identify what form the experiments will take. Chapter 14 concerns experiments and how they are reported, but prior to designing experiments the researcher must consider how they are to be used. What will the experiments show, assuming the hypothesis to be true? How will the results be different if the hypothesis is false? That is, the experiments are an evaluation of whether some hypothesised phenomenon is actually observed. Experiments involve data, code, and some kind of platform. Running of experiments requires that all three of these be obtained, and that skeptical questions be asked about them: whether the data is realistic, for example.

Experiments may also involve users. Who will they be? Is ethics clearance required? Computer scientists, accustomed to working with algorithms and proofs, are often surprised by how wide-ranging their university ethics requirements can be.

Many research activities do not have an experimental component, and instead concern principles, or fresh analysis of data, or qualitative interpretation of a case study, or a comparative reflection, or any of a wide range of other kinds of work.

However, milestones can always be identified, because (obviously, I hope!) any substantial project can be meaningfully described as a collection of smaller activities.

Two points here are worth emphasizing. First, while the components of a research project should be identified in advance, they do not necessarily have to be completed in turn. Second, we should plan research with the following attitude: what evidence must we collect to convince a skeptical reader that the results are correct? A successful research outcome rests on finding a good answer to this question.

Having identified specific goals, another purpose of research planning is to estimate the dates at which milestones should be reached. One of the axioms of research, however, is that everything takes longer than planned for.[1] A traditional research strategy is to first read the literature, then design, then analyse or implement, then test or evaluate, then write up. A more effective strategy is to overlap these stages as much as possible. You should begin the implementation, analysis, and write-up as soon as it is reasonable to do so.

For the long-term research activity of a Ph.D., there are other considerations that become significant. A typical concern in the later stages of a Ph.D. is whether enough research has yet been done, or whether additional new work needs to be undertaken. Often the best response to this question is to write the thesis. Once your thesis is more or less complete, it should be easy to assess whether further work is justified. Doing such additional work probably involves filling a well-defined gap, a task that is much better defined than that of fumbling around for further questions to investigate.

Thus, rather than working to a schedule of long-term timelines that may be unrealistic, be flexible. Adjust the work you are doing on a day-to-day basis—pruning your research goals, giving more time to the writing, addressing whatever the current bottleneck happens to be—to ensure that you are reaching overall aims.

Students and Advisors

Advisors are powerful figures in their students' lives, and every student–advisor relationship is different. Some professors at the peak of their careers still have strong views—often outrage or amazement—about their own advisors, despite many years of experience on the other side of the fence. Tales include that of the student who saw his advisor twice, once to choose a topic and once to submit; and that of the advisor who casually advised a student to "have another look at some of those famous open problems". Thankfully these are rare exceptions.

The purpose of a research program—a Ph.D., masters, or minor thesis—is for the university to provide a student with research training, while the student demonstrates the capability to undertake research from conception to write-up, including such skills as working independently and producing novel, critical insights. A side-benefit is that the student, often with the advisor, should produce some publishable research. There

[1] Even after taking this axiom into account.

are a range of approaches to advising that achieve these aims, but they are all based on the strategy of learning while doing.

Some advisors, for example, set their students problems such as verifying a proof in a published paper and seeing whether it can be applied to variants of the theorem, thus, in effect, getting the student to explore the limits at which the theorem no longer applies. Another example is to attempt to confirm someone else's results, by downloading code or by developing a fresh implementation. The difficulties encountered in such efforts are a fertile source of research questions. Other advisors immediately start their students on activities that are expected to lead to a research publication. It is in this last case that the model of advising as apprenticeship is most evident.

Typically, in the early stages the advisor specifies each small step the student should take: running a certain experiment, identifying a suitable source of data, searching the literature to resolve a particular question, or writing one small section of a proposed paper. As students mature into researchers, they become more independent, often by anticipating what their advisors will ask, while advisors gradually leave more space for their students to assert this independence. Over time, the relationship becomes one of guidance rather than management.

The trade-offs implicit in such a relationship are complex. One is the question of authorship of work the student has undertaken, as discussed in Chap. 17. Another is the degree of independence. Advisors often believe that their students are either demanding or overconfident; students, on the other hand, can feel either confined by excessive control or at a loss due to being expected to undertake tasks without assistance. The needs of students who are working more or less alone may be very different to those of students who are part of an extended research group.

An area where the advisor's expertise is critical is in scoping the project. It needs to stand sufficiently alone from other current work, yet be relevant to a group's wider activities. It should be open enough to allow innovation and freedom, yet have a good likelihood of success. It should be close enough to the advisor's core expertise to allow the advisor to verify that the work is sufficiently novel, and to verify that the appropriate literature has been thoroughly explored. The fact that an advisor finds a topic interesting does not by itself justify asking a student to work on it. Likewise, a student who is keen on a topic must consider whether competent supervision is available in that area.

Advisors can be busy people. Prepare for your meetings—bring tables of results or lists of questions, for example. Be honest; if you are trying to convince your advisor that you have completed some particular piece of work, then the work should have been done. Advisors are not fools. Saying that you have been reading for a week sounds like an excuse; and, if it is true, you probably haven't spent your time effectively.

The student–advisor relationship is not only concerned with research training, but is a means for advisors to be involved in research on a particular topic. Thus students and advisors often write papers together. At times, this can be a source of conflict, when, for example, an advisor wants a student to work on a paper while the student wants to make progress on a thesis. On the other hand, the involvement of the advisor—and the incentive for the advisor to take an active role—means that the research is undertaken as teamwork.

Over the years I have noticed that there are several characteristics that are shared by successful research students. First, they show a willingness to read widely, to explore the field broadly beyond their specific topic, to try things out, and to generally take part in the academic community. Second, they have the enthusiasm to develop their interest in some area, and then ask for advice on how that interest can be turned into a thesis project. Third, they have the ability and persistence to undertake a detailed (and even gruelling) investigation of a specific facet of a larger topic. Fourth, they take the initiative in terms of what needs to be done and how to present it, and gradually assume responsibility for all aspects of the research. Fifth, they are systematic and organized, and understand the need for rigour, discipline, stringency, quality, and high standards. Sixth, they actively reflect on habits and working practices, and seek to improve themselves and overcome their limitations and knowledge gaps. Seventh, their work *looks* plausible; it has the form and feel of high-quality published papers. Last, they have the strength to keep working despite some significant failed or unsuccessful activity; in a Ph.D., loss of months of work is not unusual.

Note that neither "brilliance" nor "genius" is in this list. Intellectual capacity is important, but many bright people do not become outstanding Ph.D. students—sometimes, because they underestimate the challenge of extended study. Indeed, I've supervised several students whose previous academic record was uninspiring but who nonetheless produced a strong thesis, in particular because they were persistent and resilient enough to pursue their work despite setbacks and obstacles.

A "Getting Started" Checklist

- Is your proposed topic clearly a research activity? Is it consistent with the aims and purposes of research?
- How is your project different from, say, software development, essay writing, or data analysis?
- In the context of your project, what are the area, topic, and research question? (How are these concepts distinct from each other?)
- Is the project of appropriate scale, with challenges that are a match to your skills and interests? Is the question narrow enough to give you confidence that the project is achievable?
- Is the project distinct from other active projects in your research group? Is it clear that the anticipated outcomes are interesting enough to justify the work?
- Is it clear what skills and contributions you bring to the project, and what will be contributed by your advisor? What skills do you need to develop?
- What resources are required and how will you obtain them?
- What are the likely obstacles to completion, or the greatest difficulties? Do you know how these will be addressed?

- Can you write down a road map, with milestones, that provides a clear path to the anticipated research outcomes?
- Do you and your advisor have an agreed method for working together, with a defined schedule of meetings?