How to write a bachelor/master thesis

Andreas Wiese

September 27, 2024

When you write your bachelor or master thesis, probably you have not written many theses like that before in your life. In fact, writing a thesis is very different from attending lectures, going to TA-classes, and doing homework exercises as you did it in the courses that you have taken before. However, many students really like working on their theses and I am sure that you will learn a lot while you do it. In this document, I list some points that can help you while you do the research for your thesis and during the writing process.

If you have any suggestions what to add or change in this document, or you have any comments or feedback in general, I am happy to hear from you: andreas.wiese@tum.de.

- Organize yourself well. When you write your thesis, suddenly there is no fixed schedule due to the courses you take (lecture every Monday from 10am-12pm, TA-class every Tuesday from 2pm-4pm, hand in your assignment sheet every Wednesday by 4pm, etc.). Instead, you need to decide completely for yourself when you work on your thesis, for how long, what you want to have finished until when; it all depends on you. Therefore, it is important that you schedule well when you work on your thesis, what you do when, etc. My experience is that for some people this is quite hard, so please do not underestimate this aspect. On the flip side, you can find a schedule that works best for you, in which you work during times of the day when you are most productive and that fits well with your other activities.
- A thesis is a marathon, not a sprint. Your thesis takes about 3-6 months (depending on whether it is a bachelor or a master thesis and on the regulations of your university). This is a long time. Therefore, I recommend you to do your thesis in the way you would run a marathon, and not the way you would run a sprint: make sure that you work regularly and steadily on your thesis, but not too much at a time. For example, it is better to work for a few hours every day than working for 14 hours one day and then not at all for the rest of the week.
- Motivation varies over time. At the beginning most students are very motivated for their thesis. Typically, after some time the initial enthusiasm degrades, and there can be moments of frustration, in particular when you do not make progress for some time. This is normal. The thing with math and theoretical computer science is that it is well possible that think about a problem for a whole day and make no progress at all. This does not feel good; neither for you nor for very experienced researchers. When you do not feel motivated, it is important that you still keep going. In such a situation, it can help that you set yourself small achievable goals for the day, so that you feel that you accomplished something by the end of the day.
- Start writing early. Even if you have all results for your thesis ready, it will probably take you quite some time to write down everything in detail, polish your proofs, write an introduction, have your chapters proof-read by friends etc. *Do not underestimate how long all this takes, for most people it takes longer than they expect.* Also, probably this is the first or the second thesis that you write which makes it even harder for you to estimate how long things take. If you have the first version of some thesis chapters ready, it is probably a good idea to send them to your advisor so that she or he can give you feedback on your writing. In this way, you can use these suggestions directly for the other chapters in your thesis. Discuss with your advisor when and how much you should send to her

or him. It is totally normal if your advisor has many comments. In particular, for this reason it is very important that you send her/him those thesis chapters quickly so that you get feedback quickly and have enough time to implement her/his suggestions.

- Write down all ideas. Ideally, write down every idea that you have, every partial result, every useful looking thought etc. on a notepad or on your computer. You will probably not include everything in your thesis that you write in this way. But the benefit is that it will allow you to, say, clear your head and it will force you to make precise what you have in mind (which is good!). And it is likely that in the process you discover aspects that you did not think of before, find potential problems with your approach early, etc. Try to write down your ideas as detailed as possible; in this way, you will discover more potential issues. Admittedly, this takes a bit of discipline, but there are a lot of benefits for you.
- Meet with your advisor regularly. In particular, meet with her/him also when you did not make progress since your last meeting, even though you tried very hard. Chances are that your advisor can give you new ideas and directions, or might even tell you that some obstacle seems to be too hard to overcome (in particular until your thesis deadline) and that you should do something else instead. Also, it helps you to keep going with your thesis.
- If you get stuck then there are several things that you can do.
 - Simplify your problem. For example, this could mean that you try to construct the smallest possible example or the simplest special case in which the problem arises that you cannot solve right now. Then focus on that smaller case and try to solve it. For example, if your want to prove some theorem about general polygons in the two-dimensional plane, you can start with the special case of convex polygons, or rectangles, or squares, or squares of unit size. Once you proved your theorem for unit squares, you can try to generalize your proof to arbitrary squares, then to rectangles, then to arbitrary convex polygons, etc. It is a bit of an art to construct the "right" special case that allows you to make progress. A good guideline is that you should define a setting in which one but not all of the difficulties of the general case arises. Another approach is to partition your problem into two smaller subproblems, e.g., one for convex polygons and one for non-convex polygons.
 - Use another approach. There are typically several different ways to approach a problem. If you get stuck with one of them, it makes sense to simply try a different one. For example, in order to compute a solution for a combinatorial problem, you can use a greedy algorithm, or a local search algorithm, or write down an LP-formulation and try to round it, or design a primal-dual algorithm for it, or formulate a dynamic program that solves it. If you tried one of these approaches for a long time and it did not work out, then it simply does not make sense anymore to keep on trying. Instead, you invest your time better if you try something else. However, this is a good topic to discuss with your advisor.
 - **Talk to your advisor.** Chances are that she or he can give you suggestons on how to overcome an obstacle, give you useful advice on whether you should try an alternative approach, or suggest a good special case to look at. Of course, you should try for some time for yourself to make progress, after all it is your thesis. But your advisor will be happy to help you if you get stuck.
- Lack of motivation. For some students, it is hard to keep going with their thesis when they lack motivation, when they are stuck, when nothing seems to work, etc. If this happens to you: don't worry, you are in good company. It is important though that you still keep going, still try more ways to solve the problem you want to solve in your thesis, etc. because, unfortunately, your thesis will not just write itself. However, for some students, this state is so severe that they effectively stop working on their thesis, completely fill their schedule with other things¹, or simply procrastinate all day. If this happens to you, I strongly advice you to seek help, not just from your thesis advisor but from some

¹Some people have amazingly clean apartments during the time they work on their thesis.

general student counselling service (most universities offer such a service). Typically, their advisors are specially trained to help you in such a situation (while your academic thesis advisor is probably not). Nevertheless, the following strategies might help you (also if you are not in such a severe situation). However, this is very individual so you need to try what works for you:

- **Regular schedule.** Work with the same schedule every day, even though you are completely flexible with your working hours. For example, make sure that you start working at 9:00am every day, then have a lunch break from 12:00pm-1:00pm, and after that continue working. In this way, you work *regularly every day* without having to invest energy into planning this. For a comparison, imagine that for two weeks straight, you go for a run every morning for half an hour. After that, you will be so much used to this that you do this automatically every morning.
- Find a work place without distractions. This can be for example a public library, a special office space in your apartment, or a room in your parents basement. The important point is that this should be a place where you go *only* for work, not for anything else, e.g., playing computer games, watching videos, etc. The idea is that for your brain it is clear that when you are there, you are there to work on your thesis, and thus it tries less to trick you into doing other things when the work becomes difficult (it will still do this at times though :-)).
- **Tell your friends.** Tell them that, e.g., from now on you will start every day to work on your thesis at 9:00am. The point is that they will ask you about this the next time you meet, and that you know that they will do this (you can even ask them to). This creates positive social pressure, because, most likely, the next time you meet you don't want to tell them that you ignored your great plans and continued procrastinating.
- **Contact with other students.** Chances are that some of your fellow students do their thesis at the same time as you, e.g., because they started studying together with you. Try to be in regular contact with them and support each other. Many people have difficult time periods while they work on their thesis, and it is good to know that you are not alone with this. Also, you could meet regularly in a public library to work on your theses together. This again creates positive social pressure, since you know that your friends will see it when you don't show up :-).
- Structure of your thesis. A typical structure for the chapters of a thesis is this:
 - 1. Introduction. You start out with a general text about the topic of your thesis. In particular, here you motivate why your topic is interesting and relevant, what applications it might have, etc. In general, you want to spark interest in the reader to read more. The beginning of this part should be written such that it can also be understood by readers who are not familiar with the area of your topic, and maybe do not know much about mathematics (imagine your friends, flatmates, family, etc.). You can begin with a very broad view and then narrow down more and more until you arrive at the specific topic of your thesis. There, you tell the reader the most relevant previous results on your topic. This is also a preparation for the next subsection in which you describe the results of your thesis.
 - **1.1. Contributions of this thesis.** You describe on a high level the results of your thesis. You can use this subsection to also give the reader an overview in which chapter one can find which result.
 - **1.2.** Other related work. If you want, you can use such a subsection in order to mention some important papers that you did not mention earlier, e.g., because they did not fit in the writing flow. You should make sure that you cite all papers that are important and relevant for your topic.
 - 2. Preliminaries. In some theses a "preliminaries" section is used to introduce basic concepts like a precise problem definition, notation, or some simple first results. However, do not put too much into such a section. Normally, it is better to introduce concepts in the places where you use them for the first time. Then, the reader does not have to go back to the preliminaries section

to recall their definition. If you can write your thesis without such a preliminaries section, this is totally fine and often a good choice.

• **3.** Chapter with results. Then you will have several chapters with the actual results of your thesis. How to structure these depends a lot on your actual results, so it is hard to give a general guideline for this.

o ...

 \circ *n*. Conclusion. Often, the last chapter in a thesis is a conclusion section. Here, you give a quick summary of your results. Try not to simply repeat things you wrote before, the reader has already read your thesis up to this point anyway. If you can, look at your results and techniques from a different angle so that you introduce a new perspective. Also, this is a good place to state interesting open problems and directions for future research.

However, depending on your thesis a different structure might be better. So please see this only as one possible option.

- Use LaTeX right. There are some things to keep in mind when you write your thesis in LaTeX.
 - **Do not use** "\\". In LaTeX, you can start a new line when you type "\\". Use this *only* if you have a *very* good reason for that. This happens very rarely :-). The idea behind this is that you tell LaTeX what you *mean*, and you let LaTeX decide things like where to start a new line or a new page. This is different from, e.g., Microsoft Word where it is common to sometimes start a new line (manually). On the other hand, it is fine that you tell LaTeX where to start a new *paragraph* since LaTeX cannot know where a new paragraph would make sense in your text. Simply put a blank line to start a new paragraph.
 - **Do not use [h] when placing floats.** When you define a float, e.g., a figure, then there is the option "[h]" to force the float to appear "here", i.e., where you define the float in your document. For example, for a figure this would be "\begin{figure}[h]". Do not use this option, unless you have a very good reason to do this (again, typically you don't :-)). The idea is again that you tell LaTeX that you want to have a float, but you let LaTeX decide where it fits best with the other content. This way, your document looks better since you leave LaTeX the freedom to place the float where it fits best.
 - Use labels. This is very basic: when you refer to something, e.g., to a lemma or a section, never write the number of the lemma hard-coded in your LaTeX code, but always use labels. For example, instead of writing "Lemma 4" you write "Lemma \ref{lem:name-of-your-lemma}". Otherwise, you need to change the numbering by hand when, e.g., you insert a new lemma before an old lemma. This is a lot of work.
- Use LaTeX efficiently. A thesis in mathematics or computer science is written in LaTeX. As a matter of fact, when you learn LaTeX there is quite a bit of a learning curve. So if you do not have much experience with LaTeX, it is useful if you start learning it soon, before you even start with your thesis. There are several programs that make it more comfortable to write LaTeX documents, like Kile, TeXnicCenter, LyX, or Overleaf. My experience is that everybody has her or his favorite programs for writing LaTeX, so you need to try and figure out what works best for you. I personally use LyX since there you see directly how your document will look like, including formulae and tables, and you do not need to worry much about LaTeX commands. For drawing figures I use IPE which has all the features you need to draw mathematical figures and it allows you to include LaTeX formulae directly in your figures (just type, e.g., \$\alpha\s) in a text box). Also, I find that it is very easy to use. Some people prefer Overleaf since you do not need to install anything on your local computer. In my experience though, it takes much longer to compile a LaTeX document in Overleaf than on your local computer (which can be annoying for a long document like a thesis). However, installing LaTeX on your local computer is not that difficult, and for your thesis I think it is well worth doing it. Under Ubuntu Linux

I recommend to simply install the package "tex-live-full" (e.g., type "sudo apt-get install texlive-full" in the console). It includes pretty much any LaTeX package that you will ever need. Under Windows you can install instead for example MiKTeX. However, figure out what works best for you, and make sure that you know LaTeX before you start with your thesis.

Also, you can check whether your university offers a template for theses, following its thesis regulations (like the template of the TU of Munich). Finally, if you have written a large part of text and believe that you will not need it at the end, instead of deleting it it might be better to comment it out or save it in some other file, just in case.

- Writing mathematics. It is important that in your thesis all proofs are mathematically precise and correct. It might take some time for you to polish your proofs until everything is spelled out precisely, every special case is taken care of, and every inaccuracy is fixed. In particular, it might take you a long time compared to the amount of text that you produce. This is normal when writing theses or papers. It can easily take you a whole morning to write a proof that is only half a page long at the end. It is important though that you invest this time: only when you write down a formal proof of your claims, you can be absolutely certain that what you believe is really correct. In many cases something seems "intuitively obviously true" but when you write down the formal proof you realize that you missed something. This happens also to very experienced researchers. On the other hand, please also give a lot of intuition to the reader. It is hard to read a mathematical text in which the formalism is precise and correct, but in which no intuition is given. For example, it is good to say something about the general structure of a section, a proof, or an algorithm, before you go into details. You may think of your thesis as a story that you tell to somebody. What structure would be good? What should come first and what should come only later?
- Web search. When you search for papers on the web, Google Scholar and DBLP are great tools. In particular, from them you can download bibtex-entries of papers so that you can cite them easily with bibtex. Also, when you do the research for your thesis topic, it can be very useful to read papers in which other people did something similar before. When you search, you might find the website of a journal etc. in which the paper was published, but you cannot download the paper without paying for it. If this happens, often you can still download the paper when you are in your university (being connected to the internet via the university WLAN) or when you connect to your university via a VPN. The reason is that universities pay for subscriptions to journals and in this way get access to such papers. If there is a freely available version of a paper, Google Scholar tries to give you a link to the PDF as you can see here (circled in red):

≡	Google Scholar	the most amazing paper ever written
•	Articles	About 910.000 results (0,36 sec)
	Any time Since 2022 Since 2021 Since 2018 Custom range Sort by relevance Sort by date Any type Review articles	propri SERVANT-LEADERSHIP: WHEREFORE ART THOU LAUREL? Laurel: (porj Servation 1) What a wonderful paper you have written here. It is marked by unusual depth of thought (porj laurelley.com L Ley - 2010 - resume.laurelley.com In this paper, I provide a snapshot of myself as a servant-leader in training in context of a beingin-training; or what feels more accurate for me is in context of a being-in-evolution. For me ☆ Save 59 Cite Related articles 20 resords The amazing paper cuttings of Hans Christian Andersen BW Brust, HC Andersen - 1994 - books.google.com dren, were so impressed by him that they remembered Andersen's paper cuttings and stories all their lives. Several of them wrote about Andersen and his amazing paper cuttings when ☆ Save 59 Cite Cited by 17 Related articles 20
	 ✓ include citations ✓ Create alert 	The greatest science-fiction story ever written EJ Stone - Nature, 2010 - nature.com "The book becomes the best book ever written for whoever collapses the wave. It's brilliant." I stared at the sheets of paper lying facedown on the printer. "You're certain I can't take just ☆ Save 59 Cite Related articles All 5 versions

However, sometimes Google Scholar is wrong here and gives you a link like the red circled ones, but this link does not send you to a free version of the paper you are looking for. Then it can help to ask Google Scholar for other versions of that paper

The greatest science-fiction story ever written	[PDF] nature.com
EJ Stone - Nature, 2010 - nature.com	
"The book becomes the best book ever written for whoever collapses the wave. It's brilliant."	
I stared at the sheets of paper lying nacedown on the printer. "You're certain I can't take just	
☆ Save 奶 Cite Related articles All 5 versions	

and if you are lucky one of them will really be freely available.

- Approximation algorithms. If you are supervised by me, chances are that the topic of your thesis is to design an approximation algorithm for some problem. For this, there is no general method or technique that always works, but in Appendix A there are some strategies that might help you.
- First the proof, then the algorithm. Again, if you are supervised by me, chances are that in your thesis you design new algorithms for some problems (e.g., approximation algorithms). Typically, in your thesis you will first describe the algorithm and then prove properties about it, e.g., that it is correct, you show a bound on its running time, etc. However, when you do the research for your thesis, it is often a good idea to do it the other way round! You *first* ask yourself how the proof could look like, and *then* you design your algorithm such that it has the properties that you need for the proof. This might sound counter-intuitive, but it is a very common way to do research on algorithms.
- Figures and pseudo-codes. It is in general a good idea to include figures that illustrate concepts or proofs. This holds in particular for figures that the readers would probably draw themselves by hand if they want to understand what you are writing. Also pseudo-codes are often useful. However, note that both exist only to *help* the reader understand what you explain. They should *never* be necessary to understand your reasoning. In particular, when you describe an algorithm, it is *not* enough to describe it in pseudo-code, but there has to be a description in the main text. The reason for this is that pseudo-codes alone are often hard to parse and to understand, simply because they do not give you the necessary space to give intuition and longer explanations.
- When to capitalize "lemma", "theorem", etc. In a scientific article, you capitalize the words "lemma", "theorem", "section", "definition", "figure", "table", "section", "chapter", etc. when you refer to a specific lemma, theorem, or section etc. by their number. For example, you write

"We will prove Lemma 5 in Section 2."

or

"Now Theorem 2.5 follows from Lemmas 2.3 and 2.4. In Chapter 4, we will present a generalization of Theorem 2.5."

However, you do not capitalize these words when you do not refer to them by their number. For example, you write

"The following lemma is difficult to prove."

or

"Throughout this section, we will define $\epsilon := 1/10$. In the following theorem, we will prove our main result."

• Mathematical notation in sentences. Please do not use mathematical notation like "<" or " \exists " in a sentence such that it takes the role of a word within the sentence. For example, do not write "We show that \forall functions f we can find a rainbow-constant which is < 10 and that $\forall d \in D : a_d \in A \land b_d \in B$."

Instead, write this sentence as "We show that for all functions f we can find a rainbow-constant which is strictly smaller than 10 and that for all $d \in D$ we have that $a_d \in A$ and $b_d \in B$." However, it is totally fine (and good) to use mathematical notation like " $d \in D$ " or expressions like "we define $f(x) := \sqrt{x}$ " within sentences.

- Useful resources
 - "Writing Mathematical Papers in English: A practical guide" by Jerzy Trzeciak and "Das ist o.B.d.A. trivial!" by Albrecht Beutelspacher (in German) which contain many helpful suggestions for writing mathematics well.
 - "The Grammar According to West" by Douglas B. West which contains many grammatical suggestions when writing mathematics.
 - "Writing a Bachelor Thesis in Computer Science" by Siegfried Nijssen which many suggestions on how to structure a thesis in computer science (and IMO most of it also applies to a thesis in mathematics)

I hope that you will write a great thesis, and that you enjoy doing the research for it and finally writing it. All the best!

Acknowledgements.

I would like to thank Holger Dell for pointing out his website on the project process of a thesis with many pointers to helpful resources, and Lars Rohwedder for pointing out the book "Writing Mathematical Papers in English: A practical guide". Also, I would like to thank Paul Deuker, Alexander Keil, Anja Kirschbaum, Alexandra Lassota, and Stefan Weltge for useful comments for this document.

References

- T. Ebenlendr, M. Krcal, and J. Sgall. Graph balancing: a special case of scheduling unrelated parallel machines. In Shang-Hua Teng, editor, SODA, pages 483–490. SIAM, 2008.
- [2] M. Garey and D. Johnson. Computers and Intractability: A Guide to the theory of NP-completeness. Freeman NY, 1979.
- [3] Jan Vondrák, Chandra Chekuri, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 783–792, 2011.

A How to design an approximation algorithm

There is no general method for designing an approximation algorithm for a problem, i.e., a method that one could simply follow and that always works (like Gaussian elemination for solving systems of linear equations). Instead, every problem is different and one needs to design an approximation algorithm specifically for it. However, there are some general approaches and techniques that are often useful. Below I list some strategies that you can use if you want to design an approximation algorithm for a given problem.

- **NP-hardness.** Before designing an approximation algorithm, it is useful to check whether your problem is actually NP-hard. For many problems this is already known. For a new problem however, this might not be known. When you try to prove that a problem is NP-hard, you can use the following approaches:
 - If the problem contains numbers in the input, like the size of an item in KNAPSACK or the processing time of a job in a scheduling problem, then in many cases you can reduce PARTITION or 3-PARTITION to the problem, and in this way prove (weak or strong) NP-hardness.
 - If the problem is about selecting a subset of some given objects according to some combinatorial constraints, sometimes it is useful to reduce INDEPENDENT SET to the problem, so that the combinatorial constraints model the input graph of INDEPENDENT SET. Note that INDEPENDENT SET is already NP-hard on planar cubic graphs, i.e., on graphs in which all vertices have a degree of three [2, problem GT20 in Appendix A1.2], which might simplify the reduction.
 - Otherwise, I would recommend you to try reducing 3-SAT to your problem by building some suitable gadgets for the variables and the clauses of a given formula. There is a special case of 3-SAT called 3-BOUNDED 3-SAT which is still NP-hard [1]. In this special case, each variable appears in at most three clauses, so you can assume w.l.o.g. that it appears either twice as a positive and once as a negative literal, or twice as a negative literal and once as a positive literal.
 - In principle, you could try to construct a reduction from any NP-hard problem; however, in my experience the problems above are the ones I would try first.
- Greedy algorithms. It is a good idea to try first whether some simple algorithmic approaches are already good enough to get a (decent) approximation ratio for your problem. Greedy algorithms are among the simplest type of algorithms. For many problems, they are not good enough to get a reasonable approximation ratio, but it makes sense to rule them out first before you try more complicated techniques. In particular, when you find counterexamples for a greedy algorithm, then you often find "difficult" instances on which you can "test" other algorithmic approaches that you try (and maybe rule out quickly that they are useful).
- Small and large objects. There are many problems in which you can partition the input items easily into small and large objects. For example, in the KNAPSACK problem you can define that an item *i* is small if its size a_i satisfies $a_i < \epsilon B$ where *B* is the capacity of the knapsack and ϵ is a fixed (small) positive constant, and you define that *i* is large otherwise. Then you can look at the small and large items separately. For small objects, greedy or LP-based methods often work well, e.g., for KNAPSACK you can get a $(1 + \epsilon)$ -approximation with both approaches. For large objects, dynamic programming or complete enumeration often works well, e.g., for Knapsack you can find the optimal solution consisting of large items only by complete enumeration in time $n^{O(1/\epsilon)}$. Often, you can show that you can round the sizes of the large items (e.g., to powers of $1 + \epsilon$), while losing only a factor of $1 + \epsilon$ in the approximation guarantee. For KNAPSACK this does not work, but for BIN-PACKING it does. Finally, you can take the best solution among the small and large items (while losing only a factor of 2 in the approximation guarantee) or try to combine your approaches for small and large items to a joint algorithm.
- **Polynomial bounded input numbers.** If your problem contains numbers, e.g., the size of an item in a BIN-PACKING instance, try first to find an approximation algorithm assuming that the input

numbers are bounded by a polynomial in the input, e.g., that they are in the range $\{1, 2, ..., n^k\}$ for some constant k. Often, if you find a good algorithm in this setting, you can sacrifice a factor of $1 + \epsilon$ in the approximation ratio in order to generalize your algorithm to the general case. Typically, this last step makes the algorithm more technical, though the key ideas are still the same as in the case of bounded input data. On the other hand, if the input data are polynomially bounded it is often easier to design a good algorithm.

- LP-approaches. A standard approach in approximation algorithms is to formulate your problem as a linear program (LP), solve it in polynomial time (e.g., via the Ellipsoid method) and then try to round it. Typically, it is easy to formulate a problem as an LP. However, rounding a given (optimal) fractional solution can be difficult. A first good step is to look for instances in which the integrality gap is large. If there is such an instance, then you can rule out quickly that this LP helps you (but maybe a stronger LP does, e.g., if you add further constraints to your LP or use configuration-LP based approaches). Often, LPs are bad for problems that have many symmetries (e.g., for scheduling a given set of jobs on identical machines to minimize the makespan) since then the LP can simply "smear" all input objects across the instance and you do not get any non-trivial bound from the LP. If you have the impression that your LP is good (i.e., has a good integrality gap) then then you can try the following techniques.
 - Rounding up all fractional variables x_i that are higher than some threshold, i.e., you set $\bar{x}_i = 1$ for your rounded solution \bar{x} if and only if $x_i \ge 1/10$. For some simple problems this is already good enough, e.g., for VERTEX COVER. Typically, this simple approach is not good enough, but it is useful to rule it out first.
 - **Randomized rounding** where the probability of rounding up a variable x_i to 1 is simply the value of x_i in the optimal solution. Together with the Chernoff bounds, this often gives a simple $O(\log n)$ -approximation which is already something. But there are more sophisticated randomized rounding schemes that give better approximation guarantees. For example, *contention resolution schemes* (sometimes also called *randomized rounding with alteration*) is also a technique that works in many settings, see e.g., [3] (as shown in this paper, the technique works even for maximizing submodular functions which is a generalization of maximizing linear functions).
 - **Primal-dual algorithms** are another idea you can try for a given LP. Often you can design a primal-dual algorithm in a straight-forward way as follows:
 - * start with the (infeasible) primal solution x = 0 (so each primal variable equals to 0) and the (feasible) dual solution y = 0
 - * raise one or all dual variables until some dual constraint becomes tight
 - * take the primal variable x_i that corresponds to the tight dual constraint, i.e., set $x_i := 1$
 - * repeat this process until the solution x is feasible
 - * do backward deletion, i.e., go through the x_i variables in the *reverse* order in which you set them to 1, and for each such variable x_i check whether you can set it to 0 without making your solution infeasible

Suppose that each primal constraint is of the form $a_i^T x \ge b_i$. Try to show the following statement: "There is a value $\alpha \ge 1$ such for each dual variable y_i with $y_i > 0$, for the corresponding primal constraint $a_i^T x^* \ge b_i$ it holds that $a_i^T x^* \le \alpha \cdot b_i$." Then you get an α -approximation algorithm.

• Guessing the value of OPT. For almost every problem, you can use a binary search framework (sometimes also called bisection search) that in each iteration gives you an estimate T for OPT. Assuming that your problem is a minimization problem, then you need only an algorithm that either asserts that OPT > T or that finds a solution with a value of at most $\alpha \cdot T$, where α is your desired approximation algorithm. Often, it is easier to design an algorithm this type than a "complete" approximation algorithm for your given problem. The binary search does not cost you anything in the approximation guarantee, and it increases the running time typically only by a logarithmic factor.

Researchers often say that this "guesses the value of OPT". Here is why: when you design and analyze your approximation algorithm, do this first while assuming that $T = \mathsf{OPT}$. Typically, if $T \neq \mathsf{OPT}$ your algorithm either gives you a solution of value at most $\alpha \cdot T$ or finds out that $T < \mathsf{OPT}$.

• **Counterexamples** can be very useful for you. If you find an example showing that some approach or technique does not work, then you can save a lot of time (which you can then invest in another approach that actually works)! Thus, if you get stuck in proving that some algorithm or algorithmic approach gives you a good approximation ratio, then try proving the opposite. In particular, if you already tried proving that some candidate algorithm is good, then you probably have some intuition for how bad instances need to look like (e.g., because for some special cases you already proved that you algorithm is good). On the other hand, if you fail at finding a counterexample, this might give you a clue for how you could prove that you algorithm actually *is* good.