Common mistakes and Good things to do when writing an NSF proposal

Things to do
Get yourself on an NSF panel. This is the single most useful way to learn how to write a proposal. Contact your program chair and volunteer, and also contact your colleagues who have been on panels before and ask them to recommend you the next time they get hit up for a panel. (I suspect, although I'm not sure, that you're more likely to get asked to be on a panel in the latter case - although NSF is always looking for people to fill panels, they tend to select from "known" or "funded" folks. But there are program managers who make an effort to get new faculty on panels.)

Visiting NSF and talking to your program chair. I'm unconvinced of the utility of this. It may depend on the program chair. It certainly won't hurt, and you can get some advice from the program chair, who will then at least know your name and what area you work in. And may be more likely to invite you to serve on a panel.

Get copies of both successful and unsuccessful proposals in your area. If someone is willing to give you copy of a proposal that was rejected, and which was later revised and accepted, and the reviews, that's the best.

GET SOMEONE TO READ YOUR PROPOSAL EARLY. Especially the summary. Ideally, get someone in your field to agree to read the summary, and then the entire proposal, and someone outside of your field to read the summary. Better to get two people to read it consecutively - you don't want to ask someone to read it too many times, partly because it's asking a lot, and partly because they'll be contaminated after the first round. But don't be shy about asking people to comment on it, especially the one page summary. Feedback is good. If you're lucky, they'll tear it apart, instead of saying "seems fine" (red ink is more useful than a pat on the head, although it is more painful).

NSF runs a CAREER workshop - I think this can be helpful. Most of the information from that workshop is here and in the slides.

Don't be afraid to trash the entire proposal and start over from scratch. Multiple times, if necessary. Or write up to the very last minute and then not submit because it doesn't feel compelling.

Make your grant office happy: get the budget/justification/facilities etc in early.

Writing style
It's OK to interleave background information/previous work with proposed work, but you have to be very clear to distinguish the two. Pick one tense/phrasing for previous work, a second for the proposed work. Separating these out with some flavor of "todo" or "I propose" at the beginning of the new work paragraph can help. Be very careful of accidentally "burying" your todo's in the middle of a paragraph.

The use of "I propose to" will get a little tedious; let it. Alternate equivalent statements "I will examine, I will experiment with, I will test", but don't get so carried away with avoiding multiple uses that you don't highlight your proposed work.

Use declaratives to discuss your existing/previous work whenever possible. Especially for women (and self-effacing folks) it is important to make it extremely

clear what your accomplishments are, what your ideas are, etc. This may make your proposal feel a bit aggressive or pushy - try to walk the line between sounding confident and arrogant. As an example:
Global warming is an on-going problem. The world needs to learn more about how global warming works and what the consequences are. Then we can combat global warming.

versus

Global warming is an on-going problem. I propose a series of experiments and models that will expose how global warming works and what the consequences are. This will lead to better ways to combat global warming.

"I" versus "we" is a personal choice; (obviously, if there's multiple PIs, it should be we). I favor "I" because, well, I'm writing the thing and I'm planning to do the work...

Don't forget that the summary page should be in 3rd person (the PI, not I). This can be extremely painful. Write it as I (or we) and try replacing all of the I's/we's with "The PI".

Proposal style and structure
As I've seen a broader selection of proposals, I've come to realize that there are two "styles" to proposals. One is hypothesis driven (here's the research question, here's how I plan to answer that question) and one is technology driven (here's the open technical problems, and my approach to solving them). Pick one or the other. I'm sure there are other styles; if anyone out there in a different field has one, please let me know and I'll add it. This is why it's important to get a copy of a successful proposal in your area - you can copy the basic flow/structure. For technology proposals you can further structure the proposal by either end-application or algorithm/technique development, whichever is a better match.

One way to maintain continuity throughout a proposal is to create a bulleted list of your main topics/ideas in the introduction, then echo that text in the subsection titles in the body of the proposal. USE THE SAME WORDS/TERMINOLOGY. Ie, if in the intro you say "Hypothesis 1: All snakes are green (explanation) " then you should have a subsection entitled "All snakes are green".

Cull, cull, cull - you don't have to get every last detail in. Better to have a few ideas presented clearly, compellingly, and well than a laundry list (as impressive as it might seem). It's ok if every last thing you've thought of doesn't make it into the proposal.

Content
Previous work
Do your homework. Especially if this is an area that's new to you. References are free, list 'em all. And, even if you don't have a lot to say about particular papers, make it clear that you've read other research group's work in that area.

It's better to pick a handful of the most relevant papers and make detailed comments on them than do a laundry list. Unlike a paper, no one expects you to comment on every paper in the area - but be sure to include citations to everything you've looked at (\nocite is your friend).

Show off what you've done in this area. Make sure to distinguish papers that are yours, e.g., we demonstrated in paper x that...

Proposed research

This is the area I can offer the least advice in, since this is very proposal-dependent. And this is, ultimately, what your proposal should be judged on. The

rest of the advice on this page is just to make sure that you've presented your ideas as clearly as possible.

If you don't have 5 solid pages of problem -> proposed solution or hypothesis ->study, you haven't got a good proposal. If all of your proposed solutions look like "we will apply technique x to problem y" or "we will investigate" without giving specifics, then you haven't got a good proposal.

For about 30% of your problems the reviewer (and you) should be 90% convinced that the proposed solutions will work. For another 45%, the reviewer should be about 60% sure. For the remaining problems, you need to convince the reviewer that these are Really Interesting problems, and even if you don't succeed in solving them, the effort will lead to New and Better things.

Absolutely new ideas versus continuing work: I've heard it said by many people that, unless you've already done most of the work in the proposal, you won't get funded. I.e., you're being funded on past accomplishments. This may be true for NIH grants, but I think it's not really true for NSF grants. I think the real problem is that, unless you've been working on a problem for a while, you won't sound very knowledgeable or specific when you talk about it. And after you've worked in an area long enough to know, at a deep level, what the issues are, you've already published several papers and... so you should expect to spend some time "proving" that your idea works before you can get it funded. The more obviously compelling/great the idea is, the less time you need to spend proving it.

Do not, do not, do not write the thesis/technique-as-hammer proposal. If you want to apply technique x to a new problem area, then you better make sure you answer the following questions:

- Why is this technique preferable to existing techniques?
- How do you propose to evaluate your technique against existing ones?
- What specifics about the application/problem make the new technique a compelling choice? (Just because it hasn't been done yet does not mean that it is worth doing...)
- What will you be able to do that no-one else has been able to do so far?
-

... think about Beta versus VHS tapes versus Laser disks versus DVDs... VHS won out over Beta format because it was just plain marketed/distributed better. It became a hard sell to convince people to adopt Beta because, although the image quality was slightly better, most people had VHS already and weren't willing to buy another machine. Laser disks came along, and although the image quality was much better, they still didn't do very well (expensive, not many disks available). Now DVD comes along and suddenly VHS has gone the way of the dinosaur. Why? The picture/sound quality isn't much better than Laser disks BUT DVDs come with all these special features...
... the moral of this story being that, unless you have a DVD of a technique, and can demonstrate it, nobody's going to buy your new technique. (And now streaming content on demand is replacing DVDs even though the quality isn't usually as good, because convenience outweighs quality…)

Evaluation
It's becoming more and more important to outline evaluation strategies for your research, especially if you're working in a well-mined area. Quantifiable measures are great things. Give some serious thought to this - how will you demonstrate that your research succeeded?
This can also make for a great outreach component - a test suite and a set of evaluation metrics.

An outside collaborator who will be using your work is an excellent way to demonstrate that your work is useful.

Equations and figures

Be very careful with equations and figures - make sure they're crystal clear (define all variables, extensive labeling/captions). I see a lot of cases where people have cut out figures/equations from their papers and stuck them in the proposal without re-writing the captions or supplying sufficient explanation. Give the figure to a grad student/fellow faculty member and ask them to tell you what's in it.

Only include equations that are relevant and necessary to explain the proposed work. Every equation should have at least a paragraph of text explaining what the equation means and how it fits into your proposed work.

That said, a clearly defined equation can elucidate what you're proposing to do. To paraphrase my advisor: "You should explain what the equations for, give the equation, then tell them why you wrote down the equation". Some reviewers are happier with text explanations, and will skip the equation, some reviewers will go straight to the equation and try to figure it out. You should accommodate both folks.

Figures: I'll say it again. Every figure should have text explaining

- What's in the figure.
- What concept/problem the figure is illustrating. I.e., why did you include this figure?
- How the figure was created (if it's an image of something).
- Explanation of any artifacts in the figure. Especially if this is a preliminary figure that isn't "perfect" yet. Don't leave the reviewer guessing why there was a big purple splotch in the middle of the image...
- Every element in the figure clearly labeled.
- Leave out elements of the figure you aren't going to explain/use
-

Tables and graphs: If you're going to show tables or comparison charts from experiments, please make sure you explain whether a low number is good, or a high number, and why. Label all axes (in legible text) and include error bars/standard deviations/p-values where appropriate.

Summaries

A timeline, or some other bulleted summary, can be very handy for the reviewers. I've seen reviewers complain about not knowing who's going to do what (i.e., what's work that will be parceled out to graduate students, undergraduates, you'll do yourself). The reviewers will (typically) have read 8-15 of these things to read, and it's hard to keep track of it all. A summary at the end they can flip to to remind themselves of what you were talking about is very helpful.

Explicitly mark the intellectual merit and broader impact paragraphs/elements in the introduction. You are now required to do this in the proposal summary - it use to be optional. But it was such a good idea that they make you do it now.

Intellectual merit: Make these precise, and only 2-4 items. Cull your intellectual merit down to the essential items - more is not better. It's very easy for those three sentences to become vague and grandiose-sounding (e.g., "Studying foo will make Great Things possible"). I've had long discussions about what is intellectual merit; I tend to think of these as being things that are technically challenging or unknown scientific questions or hypothesis that you're seeking answers to.

Broader impact: You can be a bit more grandiose here, but it still is better to have 2-3 concrete things that are believable. You can work from small to large. You must make this compelling - it use to not matter so much, but reviewers now really expect you to have thought through where your work fits in the Bigger Picture.

Another way to look at it is to make sure that each of your statements in those two paragraphs are backed up/detailed in the body of the proposal/proposal summary.

Similarly, the conclusion at the end of the proposal should summarize the proposed research, the proposed outreach, why this will benefit humanity, and why you're the one to do this. Be as specific as you can - if a sentence could be cut & pasted into another proposal in another area without change, then that sentence is too vague.

Outside collaborators
Outside collaborators, especially Real World corporations/institutions and people outside of your field, are a Good Thing. However, you need to be extremely clear about how these folks will help you, how they will be compensated, and the nature of the relationship.

If corporation/institution/person x is going to spend time and money collecting data, doing research for you, being a beta test site, etc., then this should be reflected in the budget as appropriate.

Include how feed back from your outside collaborators will be incorporated into your research pipeline. This is especially important when working with handicapped/under privileged groups. For example, if you are working with the blind on a new device for helping them get around, you need to establish that a) there's some group/person who has volunteered to test your device, and is actively interested in doing so, and preferably, that you've worked with before b) the device is addressing a problem that the blind truly want solved, not just something that you think they want solved c) there is a concrete evaluation strategy through which your blind collaborators will provide feed back.

Education/diversity/outreach
Another topic that use to be boiler plate, and is no longer. Pick one or two things that *you* want to do, are feasible for you to do, and preferably, you are already doing to some measure. Make them concrete. Better one thing you believe in than a laundry list of things you might do. Nearly every university has some outreach office now - leverage them to make your plan if you don't have one already.

Misc.
Including and excluding reviewers: The only real reason to exclude a reviewer is if you feel that person won't give you an unbiased review. Everything under review is suppose to be treated as confidential, and they make it pretty clear at the review meetings that you should "forget" everything you saw. If there's someone who's unaffiliated/conflict-free that you think would do a good job reviewing your proposal, by all means add their name. NSF program chairs are always looking for people to arm-twist into reviewing proposals.

You can always contact your program chair and ask for any advice on how to improve your proposal. They'll probably just repeat the summary back to you, but it can't hurt.

LISTEN to what your proposal reviewers said about your proposal, especially the panel summary. They are right. Well, usually. A good summary should tell you what flaws/problems to address in the next round.