UNIVERSITY OF CAMBRIDGE

A Brief Guide To Mathematical Writing

This guide is intended for first-year students, to assist them in developing a good mathematical writing style. Course information and further study skills documentation are available from the Faculty website.

Dr Gareth Wilkes

Revised October 8, 2023

Introduction

When you arrive in Cambridge, you will find that most of your lecturers and supervisors communicate and write their maths in a very particular way. For some of you this formal mathematical style may be familiar from books you have read, or perhaps from school. For others it may feel wholly new and rather arcane. The ability to write clear, precise mathematics is a new skill that you will need to pick up alongside learning the actual mathematical material of the courses themselves.

There is a great difference between a collection of sentences which happens to contain the right ingredients for a proof, and a well-ordered argument which sets out the key ideas and avoids clutter. The capacity to communicate your ideas in a clear, unambiguous and logical way is a key part of being a good mathematician. This not only helps your reader (be it a supervisor, an examiner or a future work colleague) comprehend what you're trying to say, but also clarifies your own understanding. If your response to a question from a supervisor like, 'What were you doing in this section?' is 'Um, I'm not sure. . . ' then this something you need to address!

Spending a bit of time thinking about how to write your maths well will more than pay off later. You will get more useful feedback from supervisors; you will be able to refer back to your own work more easily; and you will make fewer mistakes.

You will likely pick up a good feel for the style of formal mathematics by osmosis from your lecturers. The purpose of this brief guide is to give some specific pointers for how to write your mathematics in a clear manner, as well as to explain some of the reasons why formal mathematics is written the way it is. Some of this will feel like statements of the obvious, but sometimes the obvious needs to be stated!

A few of these tips are a case of personal preference rather than a command that you must write a certain way. You will slowly acquire your own style as you progress, and not everyone writes identically. The aim is always to make your mathematics clearly understandable, both by yourself and others. Hopefully the suggestions in this guide will help you develop such a style.

Contents

1. General tips **3**
- Rewriting your work . . . 3
- Linearity of mathematical presentation . . . 3
- If . . . then . . . 4
- Citing theorems and checking conditions . . . 4
- Diagrams . . . 4

2. Notation and definitions **4**
- Use words as well as just symbols . . . 5
- Use formatting to make finding your definitions easier . . . 6
- Choose letters that help you remember what role an object plays . . . 6
- Avoid overloading symbols . . . 6
- Define symbols only once . . . 7

3. Formatting and structuring proofs **7**
- Keep your sentences short . . . 8
- Use paragraphs to break up longer proofs . . . 8
- Lay out your claims and objectives . . . 9
- Add concluding sentences . . . 9

4. Quantifiers: $\forall$ and $\exists$. **10**
- Quantifiers in theorem statements and definitions . . . 10
- Proving a quantifier statement . . . 11
- Using a quantifier statement . . . 13

5. Glossary of symbols **16**

1. General tips

Rewriting your work

You will notice as you read through this guide that many of the tips seem to require a certain level of foreknowledge. How can you say what you're going to do in a paragraph yet if you haven't done it? How can you write a precise definition of a constant if you aren't sure what value it needs to have yet? How can you lay out the key steps of the proof in advance if you haven't finished it?

The reason for this is that writing good maths often involves re-writing your work. You should not really expect to be able to write a flawless answer to every question on the first attempt, unless it is a very short question. If you have spent a long time on a question, it is usually a good idea to step back and assess your work. What did you actually do in that argument? What were the key steps? Did you do anything which you actually never needed?

It is far better to spend a little extra time honing your mathematics, and re-writing a better proof, than to hand in what is essentially rough work. You will have a much better idea of what you actually did in your work, and you will get more useful feedback from your supervisors.

A similar tip that forms part of the rewriting process is re-reading your work perhaps a day or so after you did it. That little bit of extra distance makes it much easier to spot any flaws (both poor writing and mathematical errors) than going back over something the instant you have finished it.

Linearity of mathematical presentation

There may be many elements involved in a proof, but it is important to maintain a single flow of logic in mathematics. Your reader can only read one thing at a time, so you need to tell them which parts come before others. Having boxes scattered around a page, even with arrows between them, is strongly discouraged. Maths should generally be presented in a linear fashion.

One partial exception to this can be small side-calculations. For example, in an integration-by-parts, you may choose to write the substitutions in the right margin. This allows the reader to follow the calculations without breaking up the flow with an intervening sentence.

The guiding principle is that it should be as easy as possible for the reader to follow what you're doing.

If ... then

In usual English grammar, the 'then' part of a conditional clause is often omitted. Mathematically it is a good habit to remember to write the 'then'. If a sentence is very short, it is easy to tell which part is the condition and which part is the consequence. On the other hand, if a sentence is long, perhaps with several conditions or consequences separated by commas, then it is too unclear to simply use a comma to separate the two halves and the 'then' is vital.

Citing theorems and checking conditions

Before you apply a theorem, it is important to check that all the necessary conditions are satisfied. Failing to check conditions, and thus applying a theorem invalidly, is an easy source of mistakes.

The first time you use a theorem, it can be helpful to write it out in full yourself: the act of writing it in full forces you to consider the whole statement, not just the juicy part at the end.

Diagrams

A picture or two can greatly improve the readability of your mathematics. A diagram will rarely be a substitute for a proper proof, but can help to clarify the role of the various symbols you define.

This is most obviously helpful in a course like Dynamics where the objects have physical interpretations, but it can be of use in purer mathematics too. For example, in linear algebra it may be common for a question to begin with a blizzard of function definitions:

Let $f\colon U \to V$, $\alpha\colon V \to \mathbb{R}$, $i\colon \mathbb{R}^n \to U$, $j\colon \mathbb{R}^m \to V \dots$

A simple diagram

$$\begin{tikzcd} U \arrow[r, "f"] & V \arrow[r, "\alpha"] & \mathbb{R} \\ \mathbb{R}^n \arrow[u, "i"] & \mathbb{R}^m \arrow[u, "j"] & \end{tikzcd}$$

indicating which maps go where can be a good aid to memory, and shows at a glance which function compositions are legitimate.

2. Notation and definitions

One aspect of mathematical writing that newer students often struggle with is setting up proper notation. This is a surprisingly important part of mathematics—and one which makes the mathematics itself much more difficult when not done properly.

Let us start with a simple example of the importance of good notation. Here are two statements. One is true and one is false; can you tell which is which?

Statement 1 If $x \neq 0$ and $\alpha x = 0$ then $\alpha = 0$.

Statement 2 If $x \neq 0$ and $\alpha x = 0$ then $\alpha = 0$.

No, there isn't a typo here. These strings of symbols are identical, but their veracity depends on what kinds of objects the symbols represent. If the letters we use are not clearly defined, then the mathematics itself becomes impossible to discuss.

If I modify the statements thus, then they become much clearer:

Statement 1′ Let $x \in \mathbb{R}^n$ be a vector and let α be an $n \times n$ matrix. If $x \neq 0$ and $\alpha x = 0$ then $\alpha = 0$.

Statement 2′ Let $\alpha, x \in \mathbb{C}$. If $x \neq 0$ and $\alpha x = 0$ then $\alpha = 0$.

Now you can hopefully discern which is the true statement.

This was a somewhat silly example, but it reinforces the fact that if a symbol isn't clearly defined with its type (and any properties that it may have) stated explicitly, we do not know what operations we are allowed to perform. The simple act of writing 'let v be a vector' can help prevent common blunders like trying to divide by a vector v.

So what are some good habits when it comes to making definitions and using notation? Here are some examples.

Use words as well as just symbols

From a formal point of view, a statement 'let $x \in \mathbb{R}^3$' is a perfectly valid definition; being an element of a certain set specifies exactly what type of object you're talking about. However a more full statement like 'let $x \in \mathbb{R}^3$ be a vector', while a little unnecessary, does reinforce the nature of x.

Note that having only words and no symbols is more dangerous. 'Let x be a vector' is too vague and raises questions such as 'a vector of what dimension? And is it real or complex?'.

Including these descriptive words is not just helpful at the stage of definition. If you haven't used a symbol in a while, it can be helpful to include a reminder for your reader. Saying 'the vector $v\ldots$' rather than simply 'v' improves the readability of mathematics a surprising amount.

Use formatting to make finding your definitions easier

We may talk more about formatting later, but for now we will focus on definitions. One of the more frustrating things when reading mathematics is thinking 'Wait, what was this symbol supposed to mean?' and not being able to find it!

Burying a definition in the middle of a sentence is a good way of losing track of it. It is usually better to leave a definition in its own sentence, ideally at the start of a line.

One method that some students use instead is to put each symbol in the margin of the page where it is defined, so a quick scan up the margin will immediately locate the definition you need.

Choose letters that help you remember what role an object plays

If, in my silly example at the start of this section, I had written one statement as

Statement 1″ If $v \neq 0$ and $Av = 0$ then $A = 0$.

then many of you would implicitly have assumed that v was a vector and A was a matrix, even though these letters could in principle refer to absolutely anything. In this manner, a judicious choice of notation can take some of the burden of memory away in a complex proof involving many symbols. Keeping consistent choices of symbols, or consistent areas of the Latin or Greek alphabets, for certain kinds of object make your proofs easier to comprehend.

Certainly it makes things most confusing if you use illogical symbols! I once read the sentence 'let e be a cube and let C be an edge of e', which I'm sure you will agree is suboptimal notation...

Avoid overloading symbols

I once read the following sentence in a paper: 'let H be a Lie group with Lie algebra $\mathfrak{h}$ and let $h \in H$'. This kind of thing may sometimes be necessary, if only one choice of symbols makes logical sense, but it is rather confusing—especially if read aloud! In a similar way, some people have a tendency to use the letter 'x' to denote every unknown object in a proof, with progressively more decorations: there may be an x, an X, an x', an x_i, an $x^{(j)}$,...

Having a handful of different things labelled by the same letter is fine, especially if they form a natural set of things (the entries of a matrix, elements of a group, etc...). But don't be afraid to use a new letter to make things more readable (generally one in proximity to the old symbol, to help you remember what it is). One classic example from IA Maths is when you come to consider subsequences and wind up with a massive stack of subscripts like $x_{n_{k_{l_m}}}$. Would it not be easier to just set $r_m = n_{k_{l_m}}$, and reduce the number of subscripts by two?

Define symbols only once

Mathematics can be an unpredictable art. At the start of a proof, you do not know exactly how it will go or what properties you will need. It is fairly common to define a symbol early on in an argument and discover later that there are more constraints that need to be applied to it. In these cases, you can end up with bits of a definition scattered throughout a proof where it is difficult to find them. You may even accidentally impose contradictory conditions on an object! If you begin a proof and aren't certain what value to assign to a constant, or what subset an element belongs to, it can be helpful to leave yourself some space and come back later to emend the definition. Also, be prepared to rewrite proofs when you're done; like a piece of art, a first version is often only a draft.

3. Formatting and structuring proofs

Consider the following pieces of mathematics. Both are proofs of the same statement.

Proposition. *A function $f\colon X \to Y$ is injective if and only if $f(A \cap B) = f(A) \cap f(B)$ for all subsets $A, B \subseteq X$.*

Proof 1. If f is not injective then there are $x, y \in X$ such that $f(x) = f(y)$, so $\emptyset = f(\{x\} \cap \{y\}) \neq f(\{x\}) \cap f(\{y\}) = \{f(x)\}$. Conversely if f is injective take subsets $A, B \subseteq X$, for every $x \in A \cap B$ $f(x) \in f(A) \cap f(B)$ so $f(A \cap B) \subseteq f(A) \cap f(B)$; and if $y \in f(A) \cap f(B)$ then $y = f(a)$ for some $a \in A$ and $y = f(b)$ for some $b \in B$ but then $f(a) = y = f(b)$ so $a = b \in A \cap B$ and $y \in f(A \cap B)$, so $f(A) \cap f(B) \subseteq f(A \cap B)$ and the sets are equal as required. $\square$

Proof 2. Suppose f is not injective. Then there are $x, y \in X$ such that $f(x) = f(y)$. Then $f(\{x\} \cap \{y\}) = \emptyset$, but $f(\{x\}) \cap f(\{y\}) = \{f(x)\}$. So we have two sets $A = \{x\}$ and $B = \{y\}$ with $f(A \cap B) \neq f(A) \cap f(B)$.

Conversely, suppose f is injective and take subsets $A, B \subseteq X$. For every $x \in A \cap B$, we have $f(x) \in f(A) \cap f(B)$. Hence $f(A \cap B) \subseteq f(A) \cap f(B)$.

If $y \in f(A) \cap f(B)$ then $y = f(a)$ for some $a \in A$ and $y = f(b)$ for some $b \in B$. But then $f(a) = y = f(b)$ so $a = b \in A \cap B$. Therefore $y \in f(A \cap B)$.

Hence $f(A) \cap f(B) \subseteq f(A \cap B)$ and the sets $f(A \cap B)$ and $f(A) \cap f(B)$ are equal as required. $\square$

The mathematical content here is identical; only the formatting differs. Do you agree that the second proof is easier to understand? The way your mathematics is laid out can be just as important for its readability as the actual maths itself.

I will set out a few specific pointers below, but the key idea in all of them is to break your mathematics down into more easily digestible ideas to help your reader, whether that is a supervisor, an examiner, or simply yourself when you come back to read your work again in future.

Keep your sentences short

There can be a distinct tendency in mathematics to keep a sentence rolling on and on, as each new thought follows on from the last[1]. It greatly improves the readability of your maths if you include break points (i.e. full stops, or sometimes semicolons) to mark the end of a thought and allow the reader to absorb what just happened. You finished defining your symbols? End the sentence. You carried out one step of an argument, but now need to bring in an ingredient from elsewhere in the proof? End the sentence. You proved one direction of an equivalence, and now are moving to its converse? End the sentence. By doing so, you give your reader a space to process what you just said before moving on.

Use paragraphs to break up longer proofs

This tip is essentially a variant of the previous one. It is very rare for a long proof to consist of one single block of text. More commonly, a long proof consists of several distinct steps, and can be broken up with paragraphing. This makes your proof more digestible, and can also help you keep track of what you're doing.

There is no particular minimum length for a paragraph. If you have accomplished what you wanted to do and wish to move on to a new paragraph or start a new line, there is nothing to prevent you from doing so even if the current paragraph is very short. For instance it is common for analysis proofs to start with a paragraph containing nothing more than 'let $\epsilon > 0$.'[2]

[1] A slightly embarrassing example: when proofreading one of my own papers a few years ago, I discovered that an entire twelve line paragraph was one sentence!

[2] It is so common, in fact, that it is sometimes said that the shortest joke in mathematics is 'let $\epsilon < 0$'.

A slightly different way that some students like to mark out parts of an argument is to use indentation. If you have some smaller part of the argument that you feel gets in the way of the flow of the proof, you can indent it slightly. In this way it becomes a self-contained block with a clear start and end.

Lay out your claims and objectives

When I was preparing this guide, I asked several of my current students for their input; what advice would they give to new first years? By far the most common answer was to break down proofs into explicitly stated claims. As one of them put it to me:

> I used to write pages and pages of irrelevant stuff, but once I started listing explicit claims I started focusing on what was actually needed for the proof.

Telling the reader what you're about to do can be a great aid to understanding, whether you do it in a sentence ('we will now consider X'...) or state things in a more formal style:

Claim. ...

Proof. □

This not only helps your reader, but also makes it easier to follow the logic of your own arguments. Staring at a two-page proof with no sub-claims can be daunting, but looking at a broken-down proof means you can say 'I did this, then this, then this' and extract the key ideas from your work.

This is especially important when introducing proofs by contradiction. Setting out exactly what assumption you're making is vital for making such proofs intelligible. This is true for induction proofs as well; a simple statement 'we will prove X by induction' flags clearly what is about to happen.

A further advantage is that of cross-referencing: if there is some key claim that will be used several times in the proof, having it labelled allows you to cite it with ease. This idea also applies to equations: numbering key equations allows you to pick out important stages and refer back to them later.

Add concluding sentences

The mirror image of the previous point is the practice of ending sections with a summary statement. Without such a sentence, the reader may not notice that you completed your goal, and can sometimes expect your discussion to carry on into the next paragraph. A simple sentence to the effect of 'so now we have proved X' informs your readers that you have finished one task and are about to start another.

4. Quantifiers: $\forall$ and $\exists$.

Perhaps the most problematic symbols for new mathematicians are the two quantifier symbols: the universal quantifier $\forall$ 'for all' and the existential quantifier $\exists$ 'there exists'. This is partly because the way they are used can vary considerably depending on context. In this section we will try to break these different contexts down a little, and illustrate the correct use of these symbols.

We may begin, however, with a general pointer. Always remember that these symbols are simply abbreviations for the words 'for all' and 'there exists [...] such that'. Reading your proof aloud can often alert you to an odd use of a quantifier that may be causing problems. If you can't parse your logical statement into a full sentence, it may mean that something has gone wrong.

Quantifiers in theorem statements and definitions

The first, and arguably simplest, context in which you will meet quantifiers is in the statement of a definition or a theorem. An example, which we will be using in this guide to illustrate quantifier use, is the definition of the statement 'the sequence a_n converges to the limit l'. This definition reads:

> For every $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that for every $n \geq N$, $|a_n - l| < \epsilon$.

You may also see statements like this in fully symbolic form, which makes them more concise but less readable:

$$\forall \epsilon > 0 \; \exists N \in \mathbb{N} : \forall n \geq N, \; |a_n - l| < \epsilon \qquad (*)$$

Without thinking too hard about the actual mathematics here—that is a topic for your Numbers and Sets course—we can look at the three quantifiers in this statement and see that they play slightly different roles.

- '$\forall \epsilon > 0$': simply introducing a new symbol, which is arbitrary except for the condition that it is positive. After all, there are no other symbols yet introduced, so there is nothing on which ϵ can depend!

- '$\exists N \in \mathbb{N}$'. This N will depend in some way on ϵ: there is one N for every ϵ, but different ϵ will have different N 'belonging' to them. Some people like to write things like '$\exists N(\epsilon)$' or '$\exists N_\epsilon$' to emphasize the dependence. The notation '$\exists N(\epsilon)$' hints that you can think of these statements as being something like a function—you put an ϵ in, and get some N out.

  This is especially helpful in longer statements where there are many symbols in play: the difference between the following two statements

  $$\forall \epsilon > 0 \; \forall x \; \exists N : \forall n \geq N \; |f_n(x)| < \epsilon$$
  $$\forall \epsilon > 0 \; \exists N : \forall x \; \forall n \geq N \; |f_n(x)| < \epsilon$$

  becomes much clearer if you label the existential quantifiers to show the dependencies:

  $$\forall \epsilon > 0 \; \forall x \; \exists N_{\epsilon,x} : \forall n \geq N_{\epsilon,x} \; |f_n(x)| < \epsilon$$
  $$\forall \epsilon > 0 \; \exists N_\epsilon : \forall x \; \forall n \geq N_\epsilon \; |f_n(x)| < \epsilon$$

  In general, a symbol introduced by an existential quantifier may depend on every symbol that came before it, but will be independent of everything defined after it.

- '$\forall n \geq N$': this is part of the conclusion of the statement (generally speaking, the part after all the existential quantifiers). You can often think of the final part of such a definition as stating a property that the other symbols must satisfy. In this case, the property that N must satisfy is '$\forall n \geq N, \; |a_n - l| < \epsilon$'.

Observe how long the bullet point for $\exists$ is in this list! Existential quantifiers are often the ones that cause issues.

One more thing worth saying, both about a single statement in isolation and more generally, is that words are generally encouraged over symbols where possible. In the definition of convergence given above, the wordy version is much easier to parse than the fully symbolic string of $\forall$ and $\exists$. Your maths is written for other humans to read, and humans like to read sentences!

Let us now move on to discuss how quantifier statements like ($*$) are used in proofs.

Proving a quantifier statement

Quantifier-heavy statements often look like something from formal logic, so there is a temptation to prove them by writing a long string of formal statements connected by '$\Rightarrow$'. Occasionally this can work, but it rapidly becomes very difficult to read and comprehend. This makes it easier to make mistakes. Usually it is better to take things more slowly, and use more words.

Let us look at a proof of a simple convergence statement: that the sequence $a_n = 1/\sqrt{n}$ converges to the limit $l = 0$ in the sense defined in ($*$). The statement we need to prove is thus:

$$\forall \epsilon > 0 \; \exists N \in \mathbb{N} : \forall n \geq N, \; |a_n| < \epsilon$$

We have to begin at the start of the statement, by taking an arbitrary $\epsilon > 0$. Then we meet 'there exists N'. To show that an N exists, we must say what it is! Usually this means that we must construct it, using the number ϵ as an input.

In the case of this statement, we can easily see that $|a_n| = 1/\sqrt{n} < \epsilon$ if and only if $n > 1/\epsilon^2$, so we just need N to be some natural number bigger than $1/\epsilon^2$. Now we know what N needs to be, let us see the two ways of writing the proof mentioned above.

Proof 1: full symbolic logic.

$$\begin{aligned} a_n = 1/\sqrt{n} \quad &\Rightarrow \quad |a_n| = 1/\sqrt{n} \\ &\Rightarrow \quad \forall n \geq N \; |a_n| \leq 1/\sqrt{N} \\ &\Rightarrow \quad \forall \epsilon > 0 \; \forall n > 1/\epsilon^2, |a_n| < 1/\sqrt{1/\epsilon^2} = \epsilon \\ &\Rightarrow \quad \forall \epsilon > 0 \; \exists N : \forall n \geq N, \; |a_n| < \epsilon \quad (\text{Set } N = \lceil 1/\epsilon^2 \rceil + 1) \end{aligned}$$

□

Proof 2: more words. Let $\epsilon > 0$. Choose $N \in \mathbb{N}$ such that $N > 1/\epsilon^2$. Then for every $n \geq N$, we have

$$|a_n| = 1/\sqrt{n} \leq 1/\sqrt{N} < 1/\sqrt{1/\epsilon^2} = \epsilon. \qquad \square$$

Both of these proofs get the job done, but one of them is much easier to read and understand. I hope you will agree with me that it is the second one! The first statement, while containing nothing untrue, is very dense.

You may have noticed that Proof 2 does not include any occurrences of the symbols $\forall$ or $\exists$ in it[3]. Where did they go?

The $\forall \epsilon > 0$ statement in the desired statement became 'Let $\epsilon > 0$'. We need to prove our statement for every ϵ, but we only need to do them 'one at a time'. The 'Let $\epsilon > 0$' statement fixes one ϵ for us to work with.

Now that we have an ϵ, we can go on to show that there exists a number N depending on ϵ which has the required property. In this simple case we do this by building N as an explicit function of ϵ. The '$\exists N$' in the statement has thus become 'Take $N = \ldots$' in the proof.

[3]Except, I suppose, the place 'for every $n \geq N$' where $\forall$ is written in words.

Using a quantifier statement

Suppose we already know that a quantifier statement is true; for instance, that a_n converges to l. How might we write a proof that uses this fact? Well, the definition of 'a_n converges to l' starts with a $\forall \epsilon > 0$ statement, so for it to be of any use we need to feed it an ϵ. (After all, what use is a function if you never evaluate it anywhere?)

Perhaps we care about the value $\epsilon = 1/2$. We would naturally write the consequence of this simply as

$$\exists N \in \mathbb{N} : \forall n \geq N \ |a_n - l| < 1/2.$$

But what if we do not have such a concrete value of ϵ as $1/2$? This is where some obscurity often creeps into writing. Consider the following beginnings to a proof.

Proof 1 Let $\epsilon > 0$. Because a_n converges to l, $\forall \epsilon > 0$ $\exists N \in \mathbb{N}$ such that $\forall n \geq N$, $|a_n - l| < \epsilon \ldots$

Proof 2 Let $\epsilon > 0$. Because a_n converges to l, $\exists N \in \mathbb{N}$ such that $\forall n \geq N$, $|a_n - l| < \epsilon \ldots$

The first of these is a common phrasing found in first-year work, but it is flawed. What is the '$\forall \epsilon > 0$' doing here? It defines an arbitrary number ϵ—but we already have one! Is the second ϵ the same as the first? If it is, then we don't need to define it a second time. If it isn't, then we have two objects with the same name, which is a very bad idea.

One other thing to note about writing with quantifier-heavy statements is that if you're using several of them then you can rapidly build up a large number of different ϵs, Ns, δs and so on. You must always be careful not to define the same symbol several times; just because two different definitions use the same symbol N, they won't give you the same value of N! You can either remedy this using subscripts, or by using new letters.

To close out this section, here is an example proof where we can see how all three of these different aspects of quantifier statements interact. I have highlighted some of the quantifiers to refer to them later.

Proposition. *Suppose a_n converges to l and b_n converges to l'. Then $a_n + b_n$ converges to $l + l'$, i.e.*

$$\boxed{\forall \epsilon > 0} \ \boxed{\exists N} \in \mathbb{N} : \boxed{\forall n \geq N}, \ |(a_n + b_n) - (l + l')| < \epsilon$$

Proof. $\boxed{\text{Let } \epsilon > 0}$.

Because a_n converges to l, there exists $N_1 \in \mathbb{N}$ such that $|a_n - l| < \epsilon/2$ $\boxed{\text{for all } n \geq N_1}$. Since b_n converges to l', there exists N_2 such that $|b_n - l'| < \epsilon/2$ $\boxed{\text{for all } n \geq N_2}$.

$\boxed{\text{Let } N = \max(N_1, N_2)}$. Then for $\boxed{n \geq N}$, we have

$$|(a_n + b_n) - (l + l')| \leq |a_n - l| + |b_n - l'| < \epsilon/2 + \epsilon/2 = \epsilon$$

as required. $\square$

There are a few comments we could make here.

- I wrote out the exact definition of what '$a_n + b_n$ converges to $l + l'$' means in the statement of the proposition. This is not strictly necessary, and you would usually just expect '$a_n + b_n$ converges to $l + l'$' to stand on its own.

  Many students find it helpful to write the full statement of what they need to prove as the first step of the proof, to focus their minds on what needs to be done. I included this statement here for this reason.

- Observe that the quantifiers from the statement all appear in the proof, in the same order—these are highlighted in red boxes. There are various words in the middle, but the same flow of logic is maintained. 'Let $\epsilon > 0$' corresponds to '$\forall \epsilon > 0$', and 'let $N = \max(N_1, N_2)$' corresponds to '$\exists N$'.

  Although both statements use the word 'let', they are materially different: ϵ is arbitrary, while N has been defined exactly.

- When we use a hypothesis like 'since a_n converges to l', we did not need to put in another $\forall \epsilon > 0$ quantifier. We simply input the value $\epsilon/2$ into our 'function' and got a number N_1 with a certain property as an output.

- In a previous part of this section, I said it was poor practice to have the same symbol defined multiple times, and yet there are several '$\forall n$' quantifiers (in blue, in addition to those in red). Why is this not a concern?

  It is best to regard the two blue '$\forall n$' statements as being part of the property N_1 or N_2 is supposed to satisfy—part of the definition of N_1 or N_2. They are just throwaway dummy variables, never used again, so it doesn't cause an issue to use the same symbol for both.

If, however, we had already defined a symbol n in a more 'global' context—for example, if the proof had started 'let $\epsilon > 0$ and $n \in \mathbb{N}$'—then there would be a problem: we wouldn't know which n was being meant when we get to the later '$\forall n \geq N_1$' statements. The exact symbol used for the dummy variable when stating a property doesn't really matter; feel free to change it if you've already used that symbol elsewhere.

5. Glossary of symbols

Below is a list of a few symbols which my students have felt the need to ask the meaning of in the past.

wlog	'without loss of generality': an assumption that is made to simplify the writing of an argument without weakening the resulting theorem.
TFAE	'the following are equivalent'
QED, $\square$	Symbols used to mark the end of a proof.
$\wedge$, $\times$	Alternative symbols for the vector product/ cross product.
$-, \smallsetminus, \backslash$	Alternative symbols for the difference of two sets: $A \smallsetminus B$ is the set of elements of A that are not in B. I prefer $\smallsetminus$. A symbol that is *incorrect* here is the forward slash A/B.
$\oint$	A symbol sometimes used as a reminder that the path of integration is a loop. Does not have any mathematical difference from $\int$.
$\bar{z}$, z^*, $z^\dagger$	Alternative symbols for complex conjugation.
f^{-1}	When $f\colon A \to B$ is a function, and C is a subset of B, the symbol $f^{-1}(C)$ denotes the preimage of the set C: the set of elements $a \in A$ such that $f(a) \in C$. The fact that f^{-1} has been written *does not mean that f has an inverse function.*
$f\vert_C$, $f \restriction C$	Symbols for restriction of functions: if $f\colon A \to B$ is a function, and $C \subseteq A$, then $f\vert_C$ is the function from C to B defined by $f\vert_C(c) = f(c)$. We have simply shrunk the domain of definition to be C rather than A.
$\to$, $\mapsto$	These two similar symbols sometimes cause confusion. When defining a function, the arrow $\to$ is used between the domain and codomain: $f\colon A \to B$ means that f takes an element of A and outputs an element of B. The barred arrow $\mapsto$ is used for specific elements: $f\colon x \mapsto x^2$ is the squaring function for example.

$\mathbb{1}_B$, χ_B	Different notations for indicator functions. If B is a subset of A, the indicator function of B is the function $\mathbb{1}_B \colon A \to \{0, 1\}$ defined by $\mathbb{1}_B(a) = 1$ when $a \in B$, and $\mathbb{1}_B(a) = 0$ when $a \notin B$.
$\mathrm{GL}_n(R)$	The 'general linear group' of matrices with entries in R which are invertible, with an inverse which also has entries in R. Here R may denote various objects such as $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$ or $\mathbb{C}$.
$\mathrm{SL}_n(R)$	The 'special linear group' of matrices with entries in R which have determinant 1.