

Making a document with L^AT_EX is a little like making a computer program. You create a file with commands in it that tell the L^AT_EX processing system how to make your final document. The input file contains source code, and L^AT_EX converts this in to an output file, typically a PDF file. Here is the source code of a very simple L^AT_EX document:

```
\documentclass{article}
% Preamble here.
\begin{document}
Hello world!
\end{document}
```

Here's what each line of source code does:

1. `\documentclass{article}`
This loads a file that contains a number of definitions that set up how the document as a whole is going to look. A common choice is the `article` class. Every L^AT_EX file must start with a `\documentclass` command. I have provided a custom document class, `homework`, that you can download from the course web page, but its use is not mandatory. Note that the `\` character has a special meaning, and is used to start the name of a command (in this case the `documentclass` command).
2. `% Preamble here.`
The start of the source code is known as the preamble. You can put commands here that will be used throughout the document. Nothing in the preamble will cause any marks on the page. The preamble for this simple document has two lines, the first line with the mandatory `\documentclass` command, and a second line that has what is known as a comment. Everything on a line following a `%` sign is a comment. Comments are for the benefit of a person reading the source code, and do not result in any marks on the page. You can delete all the comments from a L^AT_EX file and it will not change how the final document appears, but it might make it harder for a human to figure out what your source code does.
3. `\begin{document}`
This signals the end of the preamble and the beginning of the document body. Everything following this point can make marks on the page.
4. `Hello world!`
Now we're writing stuff. In fact, each letter is an instruction to the L^AT_EX processor. For example, `H` means add the letter `H` in the current font to the line of text that is currently being built. L^AT_EX subsequently uses algorithms to determine where lines should be broken to build paragraphs, and how these lines should be split up to make pages. You can tell L^AT_EX to start a new paragraph by adding a blank line to your file.
5. `\end{document}`
This tells L^AT_EX that we are done describing the document and it can finish up writing the final output. Any text occurring after this point in the source code will be ignored.

L^AT_EX and mathematics

The reason L^AT_EX is used nearly universally by mathematicians is that it makes it relatively easy to typeset good looking mathematics. For example:

The side lengths of a right triangle satisfy $a^2 + b^2 = c^2$.

Here's the source code:

```
The side lengths of a right triangle satisfy $a^2+b^2=c^2$.
```

Math mode is started with a \$ character and ended with a \$ as well. Everything in between is treated as material to be typeset using the rules for math rather than the rules for regular text. Note that to get a superscript, we use the ^ character. Here's a related example:

Recall the exponent rule $x^{a+b} = x^a x^b$.

```
Recall the exponent rule $x^{a+b}=x^a x^b$.
```

To get an exponent with more than one character in it, we need to enclose the contents in a pair of curly braces. Curly braces in L^AT_EX are used to group material to be treated as a single character. But then, how do you put curly braces on the page? You use the commands \{ and \}. For example:

Let $A = \{x : f(x) = 0\}$.

```
Let $A=\{x: f(x)=0\}$.
```

All the math we've described so far is being typeset in in-line mode. This means that the math is written as part of the sentence without any extra space around it. The other common way to write mathematics is called display mode, where the text beforehand ends a line, and the math is given some extra space and centered in the middle of the page. For example:

Fermat's Theorem states that if n is an integer with $n \geq 3$, then there do not exist any non-zero integers a , b , and c such that

$$a^n + b^n = c^n.$$

```
Fermat's Theorem states that if $n$ is an integer
with $n\ge 3$, then there do not exist any non-zero integers
$a$, $b$, and $c$ such that
```

```
\[
a^n+b^n=c^n.
\]
```

The `\[` and `\]` commands are used to enclose a single line of display mathematics. Notice that the mathematics is still part of the sentence, and ends with a period. Good mathematical style dictates that everything, even the displayed math, needs to be part of a sentence. The displayed math can even be part of the middle of a sentence; you just keep going with the sentence after the `\]` command.

Frequently in displayed math, we need to break the math up into lines, perhaps with some alignment. For this, we use the `\begin{align*}`/`\end{align*}` commands.

Notice that

$$\begin{aligned}\log(4c^3) &= \log(4) + \log(c^3) \\ &= \log(2^2) + \log(c^3) \\ &= 2\log(2) + 3\log(c).\end{aligned}$$

```
Notice that
\begin{align*}
\log(4c^3) &= \log(4)+\log(c^3) \\
&= \log(2^2) + \log(c^3) \\
&= 2\log(2) + 3\log(c).
\end{align*}
```

Each line of math is terminated by a `\\` command. The `&` character is used to create alignment. The lines are typeset so that the positions of the `&`'s on each line of math all fall on the same vertical line. For alignment with binary relations, like `=` or `<` or `≤`, it is traditional to put the alignment character just before the relation. Even for displayed math split over several lines, the material all forms part of a sentence that ends with a period. It could be read out loud as “Notice that $\log(4c^3)$ is equal to $\log(4) + \log(c^3)$, which is equal to $\log(2^2) + \log(c^3)$, which is equal to $2\log(2) + 3\log(c)$ ”. It’s a bit of a run-on sentence, but still legitimate. The displayed math split into lines results in an equivalent sentence that is a lot easier to read.

The previous example also introduces the command `\log`. Many standard functions in mathematics are written in Roman font rather than in italics, and have special spacing rules. The trig functions sine and cosine are examples, as are the logarithms. Each of these functions can be typeset using an associated command: `\sin`, `\cos`, `\log`, etc.

Sometimes you want to number an equation so that you can refer to it later.

A common error is writing

$$\log(a + b) = \log(a) + \log(b). \tag{1}$$

We can see that equation (??) cannot be correct by taking $a = b = 1$. Then

$$\log(a + b) = \log(2) \approx 0.7$$

but

$$\log(a) + \log(b) = \log(1) + \log(1) = 0 + 0 = 0.$$

```

A common error is writing
\begin{equation}\label{logmistake}
  \log(a+b) = \log(a)+\log(b).
\end{equation}
We can see that equation \eqref{logmistake} cannot be correct
by taking $a=b=1$. Then
\[
\log(a+b)=\log(2)\approx 0.7
\]
but
\[
\log(a)+\log(b)=\log(1)+\log(1)=0+0=0.
\]

```

To get a displayed equation with a number, use `\begin{equation}` and `\end{equation}` rather than `\[` and `\]`. A mnemonic label `logmistake` is given in this example using the `\label` command so that you can refer to the equation number in your text without happening to know what its exact number is. You refer to the equation number using the `\eqref` command, which typesets the equation number surrounded by parentheses. If the equation number happens to change in a later revision, the text generated by `\eqref` will be kept up to date. However, L^AT_EX frequently needs to process a document twice to get all the labels up to date. There are old historical reasons for this, and it's a bit of a hassle in the modern era. Just keep in mind that if your labels look wrong, process the document again and the problems will likely go away. In some cases a third processing might be in order.

You might have wondered what the `*` in `align*` meant. It indicates that the lines in the display should not have numbers. If you use `\begin{align}/\end{align}` instead, each line in the display will be given a number. You can label a line using `\label` just as you would a normal equation. And you can indicate that a specific line should have no number using the `\nonumber` command. This business with `*` is a common convention in L^AT_EX. In fact, the commands `\[` and `\]` we saw earlier to generate a single line of un-numbered display math are shorthand for `\begin{equation*}` and `\end{equation*}`.

With a sequence of aligned equations, you sometimes want to put some justification for each step to the side of that step. For example:

Note that

$$\begin{array}{ll}
 (m+n)p = p(m+n) & \text{by commutivity} \\
 = pm + pn & \text{by distributivity} \\
 = mp + np & \text{by commutivity.}
 \end{array}$$