

Computer scientists need to learn about significant digits

I probably spend too much time reviewing research papers. It makes me cranky.

Nevertheless, one thing that has become absolutely clear to me is that computer scientists do not know about [significant digits](#).

When you write that the test took 304.03 s, you are telling me that the 0.03 s is somehow significant (otherwise, why tell me about it?). Yet it is almost certainly **insignificant**.

In computer science, you should almost never use more than two significant digits. So 304.03 s is indistinguishable from 300 s. And 33.14 MB is the same thing as 33 MB.

Why does it matter?

- Cutting down numbers to their significant digits simplifies the exposition. It is simpler to say that it took 300 s than to say that it took 304.03 s.
- Numbers expressed without significant digits often lie. Running your program does not take 304.03 s. Maybe it did this one time, but if you run it again, you will get a different number.

Please learn to express your experimental results using as few digits as you can.

[Tweet](#) 19 [g+1](#) 26

[Comments \(16\)](#)

16 Comments

1. Providing additional precision can help convince the reader that you didn't just make the number up.



Comment by A. Non — 20/4/2012 @ [9:35](#)

2. @A. Non



Because it is a lot harder to make up the number 304.03 than the number 300?

Comment by [Daniel Lemire](#) — 20/4/2012 @ [9:44](#)

3. Of course, it is 56.137% harder to make up.



Ok, now back to the serious things...

Going one step further, I would suggest replacing the numbers with charts whenever possible.

Comment by Federico — 20/4/2012 @ [9:59](#)

4. My students routinely turn in work with numbers reported to 15(!) significant digits. Drives me