

Interactive Systems Need Safety Locks

Harold Thimbleby

Future Interaction Laboratory, FIT Lab

Swansea University, Wales

harold@thimbleby.net

Abstract. *User interfaces often behave unpredictably on erroneous input — rather than blocking errors and requiring the user to correct them. The consequences of this in the context of medical devices, which may give patients undetected overdoses, can be unfortunate. The solution should include better design, including the concept of **safety locks**, that block some forms of user error.*

Keywords. Safety locks, human error, number entry, user interface design.

1. Introduction

Machine guns have *triggers*, which the user squeezes to make them fire. Machine guns also have *safety locks* so that they cannot be fired by accident. Safety locks do not significantly interfere with normal use of a gun, but they significantly reduce the probability of accidental misuse.

Why do guns have safety locks? Well, we know two things: guns are dangerous, and even the best-trained users make slips from time to time (guns may also be left unattended). Safety locks significantly reduce the chances of unintended firing.

Like a machine gun, any system controlled by a human requires some sort of trigger so that it can be made to operate when its user wants to operate it. (Of course many systems are more complex than guns: the actions triggered may require further control and refinement.)

By analogy, any system that may have unwanted effects should have its analogues to safety locks. More specifically:

- Any system that may have unwanted effects that are large compared to the final effort to make them happen, or are hard to undo should have safety locks.¹

¹ Safety locks are thus one solution to breaking the law of *commensurate effort* [8].

The only argument against safety locks is that they may add complexity or cost to a design and may in themselves cause types of error that would not or could not occur without them.

For example, safety locks on guns may slightly slow down shooters taking their first shot. Somebody may then be shot because their gun was locked in safe mode and was not immediately usable to protect them. Almost certainly it does not make sense for a safety lock to have its own safety lock, so a more considered position is that any system with potentially unwanted effects should have a safety lock (or equivalent) *unless* the safety lock itself increases the risk of problems more than the problems it is supposed to reduce. In other words, a safety lock is a design trade-off.

Even if we want to avoid safety locks so guns are faster to use, it is also a priority to ensure that a gun user doesn't get killed by their own gun! A user being killed by accident would be far more inefficient than all of the small delays of using a safety lock each time the gun is needed to be deliberately fired.

Computer systems are like machine guns. A simple action, perhaps intentional, perhaps a slip, can have enormous and quite unintended consequences.

Imagine: a single keystroke triggers sending an unfortunate email to thousands of people. A button press triggers a patient receiving an overdose of a drug. Clicking "save" overwrites a file the user wanted to keep, and there is no undo to recover all the work that has been lost.

Where are the safety locks when they are needed?

This question is wrapped up within deeper, contextual questions: sometimes the user *really* wants to do something unfortunate but which only too late they realise *was* unfortunate. The user might really want to send the email, and perhaps nothing the computer can do will make it safer, and the user only realises their error when

thousands of complaints start coming back to them much later.

Paul Cairns observes that “safety lock” is a concept more general than what is needed on obviously dangerous devices like guns. In electrical wiring systems, there are fuses. A fuse blowing is an inconvenience (it has to be replaced) but it stops an electrical fault of some sort escalating into perhaps a fire. A climber’s safety rope stops them falling uncontrollably. It does not stop them falling in the first place, and it may not save their lives (occasionally safety ropes are cut to save *other* people’s lives, so that not everyone is pulled off together), but a safety rope is a simple device that increases safety at small cost.

In this paper, whilst we use machine guns as a dramatic background motivation for safety locks, conceptually we see them as a far more general idea, and ones that could be very widely used in many sorts of applications.

Just because there are different sorts of safety lock and sometimes complicated answers to the questions does not mean we should ignore fundamental issues...

2. Ignoring safety locks in programming

A leading undergraduate textbook on algorithms [2], “the bible” of the field according to its blurb, explicitly says that it does not cover error handling. Almost 1,000 pages of text assumes all data is correct, and, if so, how to process it. The book has no sanity checks, no assertions, nothing on exception handling — in short, no safety locks. This is an example of how we train undergraduate programmers.

If any data originates from a human — or originates from a program written by a human — there is no guarantee the data is error-free. In the worst case, the program it is fed into may behave sort-of like a machine gun and have untoward consequences.

3. Ignoring safety in numbers

Number entry is one of the most basic and widely-encountered tasks.

A user presses keys, say, 1, 2, •, then 3 (where I am using • to mean a decimal point, because small dots are hard to see and confusable with periods and other English punctuation)² and then

² The symbol embossed on the key may be •. On my computer keyboard, • is combined with >, but on

the computer will convert the sequence of key presses, 12•3, to the number 12.3. Often the user will have to press ENTER or GO or some other key to indicate when they have finished entering the number.

Although we write 12•3 or 12.3 (the keys) and 12.3 (the number), to a computer these are quite different concepts, and one has to be carefully converted to the other. Humans generally think of numbers as particular sequences of digits rather than as abstract values, so the subtle differences between keys, numerals and numbers are often glossed over. Recall that computers use binary; what we write in this paper as 34, say, may be represented inside the computer as 100010 — and even then one would search in vain to find anything that *looks* like 100010, as the 1s and 0s are represented as electrical states, not as anything humanly readable. Indeed, our habitual ways of talking about numbers makes them look very easy, and correspondingly makes it quite hard to grasp how complex number entry really is.

A typical approach to converting keystrokes to numbers is as follows.

The computer starts off with the number 0. When a user hits a digit key, the number is multiplied by 10 and the value of the digit (not the key itself!) is added to it. So if the user presses the key with the label ‘3’ on it then the calculation $0 \times 10 + \text{valueof}(\text{key}3) = 3$ is performed, and then the computer has the number 3. This may seem rather obvious and perhaps a little complex for what it achieves, but exactly the same process will work on the next key too. If the user keys ‘4’ next, another digit key, then the same process is repeated: $3 \times 10 + \text{valueof}(\text{key}4) = 34$.

In this example, as can be seen, the user *keyed* ‘3’ then ‘4’ and the computer determined that they entered the numerical *value* 34. The process is repeated so long as the user is keying digits. Figure 1 shows the algorithm as it would be written in a typical programming language.

many devices, the • is on a dedicated key. Some electronics and program working together converts the location of the key on the keyboard (and whether SHIFT has been pressed, etc) to the ASCII or Unicode for a dot. Programmers rarely concern themselves with such details: in a normal program • always looks like “.” Sometimes the program code for converting keystrokes to codes is faulty, and then the user experiences so-called key bounce problems, such as ignored or extra keys seeming to be pressed.

```

var n = 0;
key = getKey();
while( isDigit(key) )
{
  n = 10*n+valueOf(key);
  key = getKey();
}
return n;

```

Figure 1. Simple code to read a number

The simple algorithm described above works very nicely until the user keys a dot or DELETE, or keys so many digits that the computer loses track of the value of the large number intended; then things get more complex and the algorithm ceases to be so simple...

We will not pursue the elementary programming details further in this paper. Nevertheless it is interesting to note that when the user has done nothing, the computer is already thinking *zero*. However, nothing and zero are not the same, and perhaps this confusion — on the computer’s part — will cause what may seem like user errors from time to time. Here’s how:

Many devices display the initial number as 0.0, which may further confuse the user, as 0.0 cannot possibly distinguish between a user who has keyed nothing, keyed zero, keyed a decimal point, keyed a decimal point followed by zero ... and so on. To illustrate the problem: if a user walks up to a device displaying 0.0 and keys 5, the display could legitimately *but unpredictably* go to any of these values: 5.0, 0.5 or 0.05. This does not seem satisfactory for devices that are used in safety critical environments, such as in healthcare.

A user may make a slip when entering numbers. Hypothetically, consider that they press 12••3. This is clearly *not* a number; it is not well-formed. What does the program do? Does it treat this as a “trigger squeeze” for some number or other, or does it recognize it as an inappropriate sequence of actions for number entry, and safely lock it out?

If the computer’s number entry algorithm is as trivial as the one sketched above, then the program will get as far as 12 and then terminate with 12.0 as the number, rather than an error.

Programmers rarely program number entry code themselves: it is needed so often it is often built-in so that there is a way of doing it automatically.

Consider the programming language JavaScript, which is one of the most widely used programming languages, as it is used to script web browsers. In JavaScript, keystrokes can be converted to numbers automatically, and few programmers worry about how this works because it is so easy to do. Yet erroneous sequences of characters are incorrectly converted. For example, 1, 2, •, •, 3 is converted to “12” — it is treated as a valid number, but is given the wrong value.

JavaScript provides several ways to convert sequences of keystrokes to numbers; for example, the built-in routine *parseFloat* converts the same string of characters 12..3 into “NaN” — now it is treated as “Not a Number.” It is then up to the programmer to check that NaN is not used as a number in the rest of the program, as this would cause knock-on problems.

Trying *parseFloat* on 1.2.3 gives 1.2. It looks like *parseFloat* reads as much of a number as it can make sense of, then returns that as the value and ignores the rest. Indeed, *parseFloat* gives 0 as the value for 0m (where ‘m’ is some non-digit, perhaps a letter) but gives NaN for m0.

In other words, JavaScript programs can try to convert a sequence of user actions into a number using any of the various built-in mechanisms provided by the designers of the language, and depending on how they do it, they will get different results.

Sometimes JavaScript recognizes errors (though it never blocks them), sometimes it ignores the error and generates some sort of number. Generally, JavaScript assumes errors can propagate, meaning that some other part of the program has to block NaN and do something sensible, rather than treating NaN as some number (maybe zero) or worse.

If the programmer chooses to handle numbers explicitly themselves, different things may happen again, especially considering the building blocks of the language (including *parseFloat*) are unreliable.

The keys of the computer or the keys of a device are the triggers that make it do things. We have shown that pressing some keys can trigger the computer to work incorrectly. Conceptually trivial, but a safety lock would only have to block invalid number entry, and like a gun’s safety lock, need not interfere at all with normal — correct — number entry.

Many more examples of the problem are given in [9], and some details of safety locks and their effectiveness are also presented (though the

present paper introduces and motivates the term ‘safety lock’ itself).

4. Complex inter-twinning

Our observations make clear that *parsing* a number, which we’ve discussed thoroughly above, *displaying* a number as the user enters it and *editing* it are very different activities, and their programming has to be very carefully integrated.

Most likely many errors happen because the overall properties of number entry are not consistent: a DELETE key may delete a key but not have a predictable effect on the parsed number. We shall see various examples below of this and other related problems.

5. Does it happen? Does it matter?

There have been no reports of mass incorrect number entry on the web. People seem to pay their bills correctly. So it *seems* like poor programming is not a serious problem.

Suppose you pay your bills online and accidentally enter an incorrect number, and you either over-pay or under-pay your bills. You probably say “oops,” and sort out the problem. It seems *you* made a mistake, so you fix the problem.

Who would think of complaining to the designers of the programs they were using, or even to the designers of the programming language?

In complete contrast, the 1994 Intel processor floating point problem was a *cause célèbre* because an unusual error, that Intel initially tried to dismiss as very unlikely, was easily reproduced by everybody who cared to try it [5]. Even though virtually nobody *needed* to do the particular sum, once people knew what the sum was, they could reproduce it and see that the result was incorrect. This undermined the credibility of the processor, and Intel had to retract and replace affected processors at cost.

In contrast, nobody has complained about number errors, and nobody has said the Intel-equivalent of “try doing $4195835 \div 3145727$ with these particular numbers as it will go wrong.” [5][6] Perhaps if somebody said “try entering $1 \cdot 2 \cdot 3$ and see what happens” would lead to improvements in the quality of programming, and force developers to introduce safety locks?

Whenever you use an interactive system that accepts numbers, try entering $1 \cdot 2 \cdot 3$ and see what

it does. Few systems, from Excel to Mathematica, from infusion pumps to web search engines, handle it correctly (see [9]).

Evidently, systems do not process number entry errors dependably. Also, users seem not to complain and not to notice. The next question, then, is: does it matter?

6. Errors that happen *and* that matter

Syed *et al* [7] report an unfortunate medical incident where a nurse entered the number 0.5 instead of 5, an error that led to respiratory arrest (i.e., potential death) of a patient. In the paper it is clearly assumed that the design of the device used was correct: the device behaved as designed, and its logs showed the nurse had entered 0.5. Problem solved; nurse error.

Nurse A apparently entered a morphine concentration of 0.5 mg per mL instead of 5 mg per mL into the infusion pump; this is what the device has logged. With the device initialized with a concentration that is ten times too low, it would naturally pump a volume ten times too much of morphine into the patient.

The patient was finally dosed with 153 mg of morphine. The then empty supply of morphine caused the device to alarm, which led to another nurse attending and detecting and correcting the error. The full details of the incident are very interesting and are beyond the scope of the present paper — for example, nurse A also made a separate error, which had the consequence of delaying the impact of the morphine overdose. Nevertheless, the patient arrested and the log of the device showed a number entry error.

The paper reports that nurse A was uncertain how to set up the infusion pump, and asked for the assistance of a second nurse, nurse B. This suggests that the human factors engineering of the device was substandard or that the training of the nurse was inadequate relative to the work they were supposed to do. While the paper does not make clear whether this was supposedly a routine job for the nurse, it does make clear that poor training for nurses was a contributory factor to the error. One might more correctly rephrase this as poor training *relative* to the complexity of the device was inadequate. Presumably, prior to the incident the complexity of the device and/or the inadequacy of the training to enable nurses to use the device effectively was not recognized by management.

It is presumably beyond the scope of the Syed *et al* paper to suggest it: but the manufacturers

simplifying and fixing their designs so they are easier and more reliable to use would be more strategic than retraining all the nurses that use the systems. Or hospital procurement should avoid purchasing systems that allow this sort of error. If procurement checked whether systems had safety locks, then eventually manufacturers would provide better systems. In the meantime, we note that it is surprising that manufacturers continue to design systems with trivial faults; the expertise to make them properly is readily available if they want to make use of it.

The user apparently made a number entry error. The infusion pump involved, an Abbott Lifecare PCS Plus II Infuser type 4100, does not have number keys, but has “increase” and “decrease” buttons. Nor does it have a decimal point. This reduces the number of keys (numbers can be entered with 2 keys rather than 11) and perhaps makes the device look simpler, but it inevitably creates modes — for example, which digit (units, tens, etc) is being increased or decreased when pressing up or down buttons? Or is a value increased and decreased, and any digit in its decimal representation might change (e.g., going from 99 to 100 in one button press changes 3 digits simultaneously)?

Does the rate of increase change the longer the relevant button is held down (typically, the first press increments by 0.1 and if held down longer increments by 1 then by 10)?

Every mode a user interface has increases the chance that the user mistakes the mode the device is in, and hence makes errors.

The paper indicates that the infusion pump was set to a *default* setting of 0.5, and it suggests (but does not make sufficiently clear) that nurse A *selected* 0.5 rather than entered it explicitly. It appears the device was set up so that 0.5 is the initial value, and the user would then increase or decrease it to the desired target value.

If so, there are various ways to explain the error, including the following:

- A single key press error could have caused the error; the keying error may not have been noticed by the nurse. We note that the device does not beep when keys are pressed, thus making it hard to keep track of which buttons work.
- Misreading the display as 5 would mean the nurse would accept it with no further action required.

Another problem (described in the paper) is the device has an out-of-range check, which should have helped protect the patient. If a dose is to take longer than 4 hours, this is blocked.

Unfortunately, entering the wrong concentration need not make the infusion of the drug take longer. It appears the device does not check concentrations (or the calculated drug delivery rates that are functions of the concentrations).

7. Other scenarios

We have discussed in some detail ways in which the nurse error may have been induced by device design. We know that other number errors cause incidents in healthcare and that the infusion pump involved in the story above is not the only style of pump. In general, then, there are other potential sources of number error.

We now briefly consider other ways the number entry error could happen (on a different device) using numeric keys rather than increase/decrease keys.

(1) We’ve already considered a simple possibility above. Assume that a • has already been pressed (perhaps some time ago, or by accident), but the display will show 0.0, which is also what it would display if • has *not* been pressed. (The standard initial display of 0.0 does not tell a user whether • has already been pressed.) If nurse A now presses 5, the value entered is 0.5 even though the nurse expects it to be 5.0 — based on the fact that on many previous occasions, pressing 5 when the device shows 0.0 has got it to 5.0.

(2) Laura Gosby has pointed out that many devices have a fixed position for the decimal point, but digits move right-to-left as they are keyed. So many devices display 0.0 as their initial display. As a user presses 123 in that order, the display would change successively through 0.1, 1.2, 12.3. (Cash machines/ATMs often work like this.) In this scenario, nurse A simply presses 5, thinks they have pressed 5, but the device would treat it as 0.5, the wrong value and out by a factor of ten.

(3) Perhaps nurse A keys 0••. This is a simple slip: 0••5 is not the start of a well-formed number. Nurse A recognizes this error, and presses the DELETE or CANCEL key to correct the erroneous keystroke. On any sensible device (like Microsoft Word) pressing •• DELETE has the meaning same as • alone.

Written around 1790BC, the Code of Hammurabi says things like, “If a builder builds a house for someone, and does not build it properly, and the house falls down and kills its owner, then the builder shall be put to death.” Likewise, today’s programmers should be plugged into their devices: not only is it a deterrent for bad programming, but if the programmer is killed by their own bad programming, they won’t have any children. *Eventually* evolution will take care of improving programming standards...

Figure 2. Updating the Code of Hammurabi?

However, most medical devices do not keep track of how many decimal points a nurse has keyed: a number either has no decimal points or one. So 0•• is recorded by a device as 0• with only a single decimal point. But then nurse A presses DELETE to correct the error. The result will be as if the nurse pressed 0, not 0•, because effectively both dots keyed by the user have gone. Next, nurse A presses 5, and the device has got 05 as the number entered, presumably equal to the numeric value 5, yet nurse A believes they have entered 0.5.

This is a tenfold error (but, as it happens, in the wrong direction for the specific Syed *et al* scenario). Nevertheless it illustrates how a user’s reasonable expectations (learned from other familiar systems like word processors) may be seriously undermined by poor programming.

(4) Another possibility is that the device initially displays 0.0, and as the user keys in digits and dots, the display updates. The nurse keys 5, and the display shows 0.5; if the nurse continued and keyed 0, the display would update to 5.0, as the 5 “scrolls” to the left. Who knows? But this is a plausible way that nurse A thinks they have keyed 5, but the device logs 0.5.

(5) The nurse keyed • by mistake as their first key press, and noticed this slip. The nurse hit DELETE but the delete did not work, perhaps because the nurse did not press the key hard enough. Keying the digit 5 next would have led to the device treating the number as 0.5, not as 5.

(6) Another, sadly common, possibility is that the device has timeouts. The nurse presses 0• and then is maybe distracted for a few seconds (perhaps to say something or attend to the patient) then continues by pressing 5. The nurse knows they pressed 0•5, but the device timed-out after the • and ignored it. The nurse has, so far as the device is concerned, entered 5.0.

(7) On devices with up/down keys, the up/down keys may change the value displayed by 0.1. The user pressed UP and the number being entered cycles via 0.9 to 0.0 (or 0.1), because some other operation is required for UP to increment units and tens digits. On seeing the number reach 5, the nurse may have thought “I have increased the number, and it has changed to 5, therefore I have finished.” Unfortunately, this is also a misreading error: the number is 0.5, even though the nurse *saw* it increase from 0.5 to 0.6, 0.7... etc. This error would be even more likely if the user interface accelerates the rate of change the longer UP is pressed.

(8) We know that the infusion pump involved uses 7 segment digit displays. A display of 0.5 (which was selected) rather than 5 (which should have been selected) may easily have been misread, as decimal points in 7 segment displays are not very salient. Possibly the decimal point was faulty; possibly the machine was positioned so the dot could not be seen because of parallax. Possibly 0.5 is displayed as 05, with a slightly smaller decimal 5 but no decimal point at all.

(9) The hypothetical scenarios above were limited to the specific 0.5 or 5.0 confusion, though they obviously generalize to other values, such as 0.1 and 1.0 confusions. The literature explores many other forms of numerical confusion — for example, Johnson *et al* [3] show how keying 135•0 can be taken as 1,350 because the device gratuitously ignores decimal points on values larger than 100. Cairns and Thimbleby [8] show how keying 10•5 gets 10.5 but 100•5 gets 1,005 on the Baxter Colleague. In these two examples, the device’s handling of • depends on contingent modes that may not be, and probably are not, obvious to the user.

(10) All the hypothetical examples above assume the user made a slip in the operation of the device. The solution to these sorts of problem are better device design and/or better user training. (We prefer better device design, since users will eventually make slips regardless of their training.) A different possibility is that the user *intended* to enter 0.5, but by mistake. Perhaps the prescription was misread? Perhaps there was a smudge before the 5 that was misread as a •? Perhaps if the dose had been written as 5•0 it would have been less likely to have been misread as 0•5? Most probably, if the training of everybody was that . is *never* written when • is intended, the (hypothetical) misreading would be far less likely.³

³ A user may intend to make a mistake, for example knowing perfectly well what the correct dose is; this is then a *violation*.

Sadly, these ten hypothetical scenarios (in addition to the ones possible on the actual device involved in the incident, which we described earlier) are all plausible: user interfaces are often badly designed, and it is all too easy to imagine the errors not being blocked by safety locks.

In summary, there are many hypothetical ways that a device can log an apparent “user error” which has been induced by design, not by unreasonable behavior on the part of the user. Logs should at least record the exact keystrokes the user performed and their timings, not the final value the device somehow works out, perhaps erroneously, from the keystrokes.

- ***Device logs should allow investigators to distinguish between different possible causes of error.***

One can infer in the Syed *et al* case that the device log was *not* sufficiently clear to suggest alternatives other than “obvious nurse error” to the authors of the paper.

8. Discussion

The list of possible causes of error is not intended to be exhaustive, nor are all possibilities suggested equally likely on the particular device at the centre of the reported case. The point is to raise a variety of ways in which the absence of safety locks can cause problems. Almost certainly one or more safety locks on the device in question failed or, more likely, were never designed into the device.

Sadly, without more detailed information (of the device condition and set up, including default values, and of the incident itself) one can only speculate as to the specifics.

The Syed *et al* paper mentions documented cases of concentration errors made before on the same device. Evidently, concentration errors are well known and encountered relatively often. One wonders whether there is a systematic reason for this class of error other than independent human error. The unvarying common factor in all the errors is the device (and its design).

The paper says,

“While the primary error involved incorrect programming⁴ of the PCA⁵ pump.”

We would beg to disagree. The primary error, in our opinion, lies in the design of the device. Moreover, fixing the design would solve many of the paper’s recommendations, such as improved training. Why have better training for something that is over-complex, as it would be more strategic to improve the device — and reprogramming the device (e.g., in a firmware update) could be done without any training costs.

The Syed *et al* paper makes six recommendations, including better nurse training (they do not mention that perhaps prescriptions should be written out more carefully, which would be an issue of pharmacy or consultant training).

None of their recommendations cover better procurement of devices (why are bad devices purchased?) nor better design of devices (why are bad devices designed in the first place?) Nor is there any explicit learning from the investigation that might lead to better device design.

Although the paper was written from a medical perspective, from our point of view, like many if not all such papers it is too vague about the design and use of the device. The assumption is that the nurse made the error, not that the device induced or contributed to the error.

9. Conclusions

Syed *et al* is only one paper, but its concerns and approach are not unusual [1]. There was a medical mishap, and the paper makes it clear that this error was not a one-off event for this device. The healthcare implications are discussed. Maybe there will be some management changes. Design was not criticized in detail, except tangentially in mentioning a human factors paper that shows that human factors can reduce error [4]. The paper leaves unexplored whether human factors redesign would have avoided the problem discussed.

In summary:

- Many devices have no safety locks in number entry;
- Number entry errors may cause significant problems;

⁴ **Programming:** i.e., the sequence of keystrokes performed by the nurse, *not* programming as in the software programming of the device.

⁵ **PCA:** patient controlled analgesia; i.e., an infusion pump with features for the patient to have some control over their drug dose.

- Analyses do not discuss details of use in sufficient detail to be certain of relevant design factors;
- Although human factors can improve performance and safety, no connection is made between particular incidents and general solutions.

Although the present paper has concentrated on number entry (because of its familiarity, clarity and ubiquity), there are problems with all forms of user interface. Number entry is, for the sake of exposition, merely easier to demonstrate as flawed.

Machine guns are dangerous, but they are made considerably safer without compromising their effectiveness by having safety locks. Now imagine this: there is a machine gun available on the international market *without* a safety lock. I think we would all know about it: films would be made of exploits with it!

Now imagine that somebody has been shot by a machine gun, but nobody mentions that the machine gun involved has no safety lock. If a gun has no safety lock, it would be likely to be mentioned in any accident investigation. Yet nobody mentions it. For some reason, people are not thinking about it.

Bizarre with guns, maybe, but the story implies nobody is familiar with safety locks when it comes to programming and complex devices like hospital infusion pumps. Nobody notices safety lock absence. This paper has argued, as is more obvious with small arms, that safety locks in programming could lead to lives being saved. Elsewhere [9], we've shown they could halve mortality.

The need for safety locks on machine guns can be easily and dramatically demonstrated. Imagine there is no safety lock: press the trigger, and the gun sprays bullets ...and the gun's recoil creates chaos...

Without a safety lock, a gun is *clearly* lethal. Guns have the advantage, both for the rhetoric of this paper and in reality, that there is a very small psychological distance between pressing the trigger and a dangerous bullet emerging noisily!

With any other device, try entering 1•2•3 (or make some other keying error) and see what happens. Sadly, patients with an avoidable overdose of morphine die quietly without any noisy drama.

10. Acknowledgements

This work was funded by EPSRC grant nos. EP/G003971/1 and EP/G059063/1. The author is grateful for discussions with Chitra Acharya, Paul Cairns, and Laura Gosby.

11. References

- [1] Blandford, A, Buchanan, G, Furniss, D, Curzon, P, Thimbleby, H. Few are Looking: Invisible Problems with Interactive Medical Devices, *Workshop on Interactive Systems in Healthcare (WISH)* in press, Proceedings ACM CHI 2010, Atlanta.
- [2] Cormen, TH, Leiserson, CE, Rivest, RL, Stein, C. *Introduction to Algorithms*. 3rd edition. MIT Press, 2010.
- [3] Johnson, TR, Tang, X, Graham, MJ, Brixey, J, Turley, JP, Zhang, J, Keselman, A, Patel, VL. Attitudes toward Medical Device Use Errors and the Prevention of Adverse Events, *The Joint Commission Journal on Quality and Patient Safety*, **33**(11):689-694, 2007.
- [4] Lin, L, Vicente KJ, Doyle DJ. Patient Safety, Potential Adverse Drug Events, and Medical Device Design: A Human Factors Engineering Approach. *Journal Biomed Inform* 2001, **34**: 274-284.
- [5] Moler, CB. A Tale of Two Numbers, *MATLAB News and Notes* 1995, 10-12.
- [6] Soderquist, P, Leeser, M. Area and Performance Tradeoffs in Floating Point Divide and Square-Root Implementations, *ACM Computing Surveys* 1996, **28**(3):519-564.
- [7] Syed, S, Paul, JE, Hueftlein, M, Kampf, M, McLean, RF. Morphine Overdose from Error Propagation on an Acute Pain Service. *Canadian Journal of Anesthesia* 2006, **53**(6):586-590.
- [8] Thimbleby, H. *User Interface Design*. Addison-Wesley, 1990.
- [9] Thimbleby, H, Cairns P. Reducing Number Entry Errors: Solving a Widespread, Serious Problem. *Journal Royal Society Interface*, in press, 2010.