

A New Informatics Geography

E. Coiera

Australian Institute of Health Innovation, Macquarie University, Australia

Summary

Introduction: Anyone with knowledge of information systems has experienced frustration when it comes to system implementation or use. Unanticipated challenges arise frequently and unanticipated consequences may follow.

Objective: Working from first principles, to understand why information technology (IT) is often challenging, identify which IT endeavors are more likely to succeed, and predict the best role that technology can play in different tasks and settings.

Results: The fundamental purpose of IT is to enhance our ability to undertake tasks, supplying new information that changes what we decide and ultimately what occurs in the world. The value of this information (VOI) can be calculated at different stages of the decision-making process and will vary depending on how technology is used. We can imagine a task space that describes the relative benefits of task completion by humans or computers and that contains specific areas where humans or computers are superior. There is a third area where neither is strong and a final joint workspace where humans and computers working in partnership produce the best results.

Conclusion: By understanding that information has value and that VOI can be quantified, we can make decisions about how best to support the work we do. Evaluation of the expected utility of task completion by humans or computers should allow us to decide whether solutions should depend on technology, humans, or a partnership between the two.

Keywords

Human-computer interaction, utility, expected value, value of information, workaround, communication

Yearb Med Inform 2016:251-5

<http://dx.doi.org/10.15265/Y-2016-018>

Published online November 10, 2016

The journey from information technology design to use is not a forgiving one. No one doubts that health information technology (IT) is essential to the operation of modern, safe, and effective health systems. Yet our experience is peppered with frustration and the many unanticipated consequences of IT [1]. Clearly we still do not understand enough about the role that technology should play in the organisations within which we work.

Without a compass for the journey from design to implementation, one can travel through some very strange territories indeed (Figure 1). Unanticipated challenges, ranging from human-factors problems to user workarounds can subvert the intent of technology. However we should be able to conceive other routes and other geographies to explore. In this paper, I look at a few foundational informatics concepts, which, if understood deeply, can help guide us to the successful implementation of health information systems.

What is the Fundamental Purpose of Health Informatics?

“Medical informatics is as much about computers as cardiology is about stethoscopes. For those who have studied the application of information technologies in medicine, the last decade has delivered one unassailable lesson. Any attempt to use information technology will fail dramatically when the motivation is the application of technology for its own sake rather than the solution of clinical problems. [2]”

When I wrote those words in 1995, I was making a statement about the applied nature of informatics and the pre-eminence of outcome over method. Even then, it was commonplace to see clinical information systems fail because they were first and foremost conceived of as technology projects, designed in a vacuum devoid of clinical experience.

Friedman’s “Fundamental Theorem” of Informatics says something very similar [3]. If we take a human (H) and a computerized information system (C) the theorem states that:

$$H + C > H$$

The fundamental theorem does not imply that using technology is always better than not. It is a statement of intent. The use of technology *should* leave us better off than not using it.

So, at its very simplest, the purpose of informatics is to harness technology to make us better at executing tasks than if we were unaided. How it does this is not so simple however.

Information as Universal Panacea

An unstated assumption in informatics is that *adding information to a human process using technology can only improve it*. The bigger the analytics and the data are, the better. Yet the empirical evidence reveals that while IT sometimes improves clinical processes, often it makes no apparent difference and sometimes can make things worse. The research evidence tells us that [4]:

- Electronic Health Records (EHRs) decrease nurse data entry time but increases

the time for doctors. EHRs can improve record completeness but are not associated with improvements in the quality of care (e.g. length of hospital stays, death rates).

- Care pathways and plans can reduce practice variation and increase compliance with standards of care. They can improve process outcomes (e.g. appropriate test ordering through the use of drug order sets) but typically they have no impact on clinical outcomes.
- Telehealth interventions increase patient satisfaction and improve patient outcomes in some cases (e.g. in chronic care). Surprisingly however, in many cases telehealth is not cost-effective.
- Decision support systems typically improve the safety and efficiency of care and also improve patient outcomes, but are not widely used.

These are not the results we were looking for. There is often disappointment when IT evaluations produce such negative results. Technology advocates often make excuses that such results are not generalizable because “Things are different in our organisation”. Detractors will argue in return that the informatics venture is fundamentally flawed. So, how do we explain these perplexing negative results? Informatics theory can help. Rather than being a surprise, these outcomes are at least explicable and were probably predictable from the start.

If adding information to a process leads to improvement, then by definition that information must bring added value. One way of quantifying that value is to ask how many times information (such as a patient’s clinical notes, or a treatment guideline) must be read before there is a measurable change in outcomes. Metrics such as the *number needed to read* [5] and the *number needed to benefit from information* [6] attempt to correlate access to information with any resulting benefit. A clear implication of these measures is that some information exposures lead to quantifiable change but not all do.

The value of a change in outcome due to new information can be calculated using classic decision theory. For example, a diagnostic test might allow a patient to avoid a risky treatment and instead receive a less risky but equally beneficial one. The value

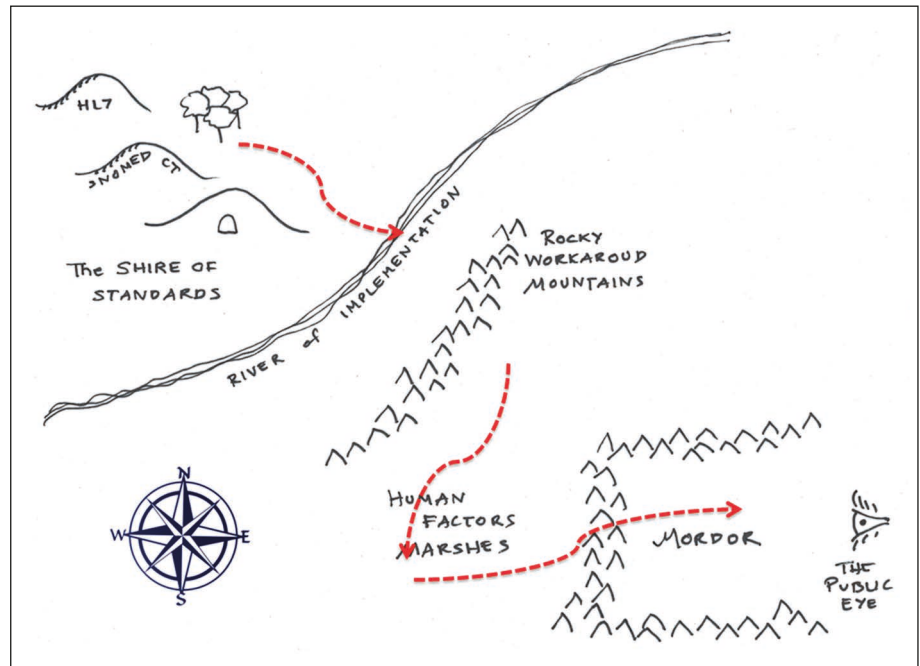


Fig. 1 Today’s informatics geography contains many perils for those who take the journey from system design to routine use. That journey has many parallels with the experiences of a travelling Hobbit in Tolkien’s *Lord of the Rings* [14]. Many of us spend all our days in the Shire, which is a wonderful place. Here we conceive new information frameworks, architectures, terminologies, and ontologies, intended we are sure for widespread use throughout the land. We never however leave to find out if they are. Everything however, changes when you leave the Shire and cross the River of Implementation, and bring a real information system into actual use. The people you meet bring you unanticipated problems. When you cross the Workaround Mountains you meet folk that are frustratingly expert at doing what they want to do, no matter what your technology tells them to do. The people in the Human-factors Marshes are expert at doing exactly what your technology tells them to do, even when it is the wrong thing to do. If you are really unlucky your journey will take you to the very dark place of Mordor, the home of large-scale IT failures. When you are here, you are always under the watchful, unforgiving, gaze of the Great Eye of Public Opinion. (Figure loosely adapted from [14]).

of the information from the test is based on quantifying avoided risks or added benefits. Specifically, the *Value Of Information* (VOI) is the *difference* between the value of persisting with the present state of affairs and the value of embarking on a different course because of new information [4]. VOI is zero whenever new data do not change outcomes. VOI is derived by first calculating the *expected utility* (EU) of the two outcomes, which is simply the *likelihood* of each event multiplied by its *utility*. The expected utility of a treatment is the probability of survival with the treatment multiplied by a utility value which estimates the cost, pain, and suffering of undergoing the treatment. VOI is just the difference in the EU of any two competing options [7 8] i.e.:

$$\text{VOI} = \text{EU} (\text{Option 1}) - \text{EU} (\text{Option 2})$$

A long *information value chain* connects the act of seeking information, through

to making a decision and that decision having an impact (Figure 2). At each step in the chain there is potential for a ‘loss’. Not every input (e.g. new information) generates an output (e.g. a modification of a decision). New information does not always lead to a decision being changed, only some changed decisions will result in a change in the process of care, and only some process changes will actually have an impact on outcomes. This means that the number of events is typically higher earlier in the chain. Equally the value of individual events appears to increase further down the chain (compare the value of reading a record to the value of undertaking a different treatment) [4].

We can calculate the EU at each step of the chain, and a VOI when comparing different approaches to supporting that step. For example, when seeking information in a patient record, the probability of receiving

new information depends in part on the completeness of the record, and utility might be based upon the speed and difficulty of information retrieval.

The expected utility from system interactions through to patient outcomes should have quite different profiles for different classes of information systems (Figure 3). For example, a telecare system will probably

maximize EU at the early doctor-patient interaction stage by avoiding the cost of face-to-face interactions. As telecare is often a substitute to normal care, there may be no expectation of changing clinical outcomes – simply a desire to provide the same quality of care at distance. Process improvement goals are reserved for other functions such as test ordering systems, care pathways, and clinical

guidelines. A decision support system in contrast is designed specifically to improve decision-making and outcomes.

With this perspective it is not surprising that an EHR does not improve clinical outcomes. Rather than seeing this as a negative result, it is exactly what we would expect since accessing records will have impact earlier in the value chain. The EHR is too upstream from decision making to easily demonstrate significant clinical changes. Information can only create value if it leads to something new and beneficial happening in the world - data alone do nothing. If we do want to see such changes, we would need to couple the EHR with clinical pathways to optimise process outcomes, or with decision support to alter decisions and clinical outcomes.

The Substitutability of Human Work with Computer Work

Another often unstated assumption in informatics is that all *human tasks can be directly replaced by automation*. Although many tasks can be automated, it is not always sensible to do so. Many tasks are fluid, evolving, and require high levels of interaction between humans to be executed. Sometimes the details of a task only become clear at execution time. Sometimes a task is about sense-making, where information is shared to allow multiple individuals to reach consensus. Such tasks typically occur in the *communication space* and are probably better supported by informal tools like communication technologies rather than by formal computational tools that require explicit and a priori task description [9].

Misunderstanding what should be automated leads to surprises. After an information system is embedded in an organization, it may not be used as often as expected, or it may lead to unintended consequences. Sometimes users subverted the system's original intent by creating workarounds.

Such workarounds happen for many reasons. Firstly, system designers may confuse

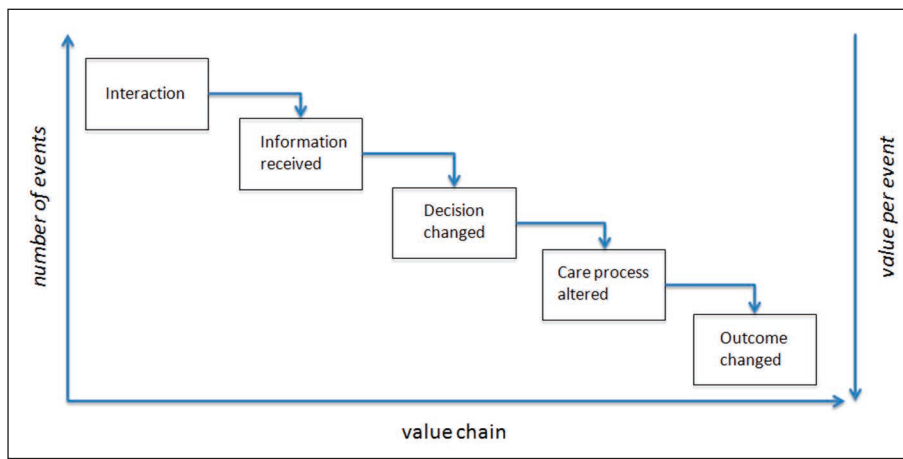


Fig. 2 For any decision task, there is an information value chain that starts with a user interacting with an information source, and goes through many steps before outcomes are observed in the world. The number of events is typically higher earlier in the chain, and the value of events is higher further down the chain. Combining event frequency (or probability) with event value (or utility) provides the expected utility of each point in the chain. (From Guide to Health Informatics (3rd Ed.) [4])

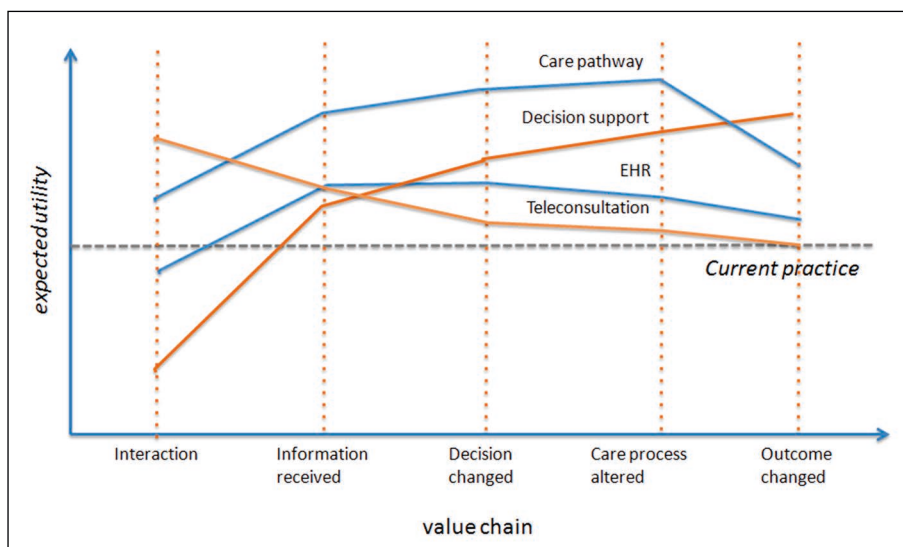


Fig. 3 The expected utility for a technology intervention may vary at any step in the information value chain. This figure illustrates hypothetical expected utility profiles for four different classes of technology interventions compared to a common non-technological baseline. An intervention (i) may improve the quality of interactions in a health service but provide little additional information compared to current practice (e.g. teleconsultation); (ii) may optimize the quality of information capture (e.g. Electronic Health Record); (iii) may improve the quality and efficiency of clinical processes (e.g. electronic care pathways) or (iv) may intervene directly in the decision-making process to improve clinical outcomes (e.g. decision support systems). Some portions of the profile may dip, and show a net cost rather than benefit (e.g. interacting with EHRs requires more time than normal for doctors). (From Guide to Health Informatics (3rd Ed.) [4])

‘work as imagined’ with ‘work as done’ [10]. What people say they do and what they actually do are not the same. As a result, digital task descriptions may ignore the necessary variations needed to adapt task execution to local context. For example, IT systems often ignore the cognitive realities of their users. High levels of cognitive load, external interruptions while the system is being used, or multitasking which requires users to split their attention can all interfere with the execution of ‘work as imagined’ encoded in automation [11].

‘Work as done’ is thus often non-linear. Pathways and dependencies between tasks may be hidden until an unexpected work context reveals them. Only then do we realise that the task as described in the computer system is inaccurate. Work is also adaptive to changing needs over time. As goals change users may need to create workarounds by taking existing tools and modify or re-purpose them. This means that in complex adaptive settings computerised work processes can be brittle unless there is a simple way to flexibly adapt the technology. Communication tools often become the bridge between work as imagined for an IT solution and work as done.

We should thus see workarounds as gifts. Rather than representing a problem with the way users engage with a technology, workarounds are clear signals that there is a mismatch between work as imagined and work as done. Indeed, we can think of *workarounds as repairs*, providing *missing information, new pathways or tools* to improve a system’s fitness for purpose. They are users’ attempts to fix inadequacy in design and to meet emergent or unanticipated needs [12].

It is thus a mistake to focus only on what computers can do. We also need to see where humans excel. The work environment is a *partnership* between a user and technology. The weight of the partnership shifts from one to the other depending on the circumstances. To recognize the important contribution that humans bring to the partnership, we can extend the Fundamental theorem to a pair of goals:

$$H + C > H \text{ and } H + C > C$$

Which Tasks Should We Computerize?

If not every task is suited to the application of technology, the next question we should ask is “Which tasks should we computerize?” [13]. Knowing the answer would provide a powerful compass to guide system designers, builders and users.

Our territory is the universe of tasks that we need to undertake, and the question we need to answer is which tasks are best completed by machine, human, or a partnership between the two. We have seen that the value of completing a task can be assigned an expected utility. The task space is created by plotting the expected utility of task completion by computer (C) or by human (H) (Figure 4). The equivalent benefit line occurs where $EU(C) = EU(H)$. Above that line there is greater utility in using automation. Below the line, outcomes are better when the task is completed by humans.

The task space divides into four quadrants (Figure 5). The top left quadrant, where $EU(C) > EU(H)$ and $EU(H) < 0$, is the natural space for automation. Tasks in this *automation quadrant* are poorly executed by humans and easily undertaken by computers e.g. completing large numbers of calculations in very short time frames. Calculating the location and dosing for radiotherapy is one example.

The bottom right quadrant where $EU(H) > EU(C)$ and $EU(C) < 0$ is the *communication space* where humans excel. These tasks might be poorly defined, dynamic in nature, perhaps evolving, and are ones in which humans are adept [9].

The bottom left quadrant is a task execution Badlands because neither humans nor computer are strong here. If forced to work here by necessity, then working above the equivalence line in a ‘mitigation zone’ aided by technology gives the ‘least worst’ outcome. Below the line is a dead zone where outcomes are invariably poor and we cannot be helped by technology.

It is in the top right quadrant where effective partnership between humans and computers occurs. Outcomes of task execution are always positive in this space. Above the equivalence line is the territory captured

in the ‘Fundamental Theorem’ by Friedman where the use of technology improves the unaided human. Below the equivalence line is where humans are better at a task than computers but where computers still add value. It is also the place where we are more likely to see workarounds as humans re-purpose technology beyond its intended design.

Before we set out to automate human tasks, we should ask ourselves which quadrant we are in. Understanding the answer has fundamental implications for the strategy taken and the likelihood of success (Figure 6). We can now also, maybe cheekily, provide a final reformulation of the Fundamental Theorem. Thinking not just about the value of information but also considering the value of automation or the value of a human as information processor suggest that:

$$\text{Value of human} = EU(H) - EU(C)$$

$$\text{Value of automation} = EU(C) - EU(H)$$

Conclusion

This paper began by asking what the fundamental purpose of informatics was. The simple and obvious answer is that informatics exists to make us better at our tasks than if we were unaided. However, we have also been reminded that computerization does not always lead to a better outcome and that phenomena such as unintended consequences and user workarounds can cloud the picture of technology benefits.

By understanding that information has value and that the value of using or not using technology can be quantified, we can make decisions about how best to support the work that we do. Asking explicit questions about the best way to support tasks helps to decide whether solutions focus on the technology, the human, or a partnership between the two.

Information value does not tell us which steps to take in solving problems nor whether any individual journey will be hard or easy. It tells us much about the kind of partnership we should create between humans and automation. That is no small thing. If you know where you are, then you have a better of chance of getting to where you want to go.

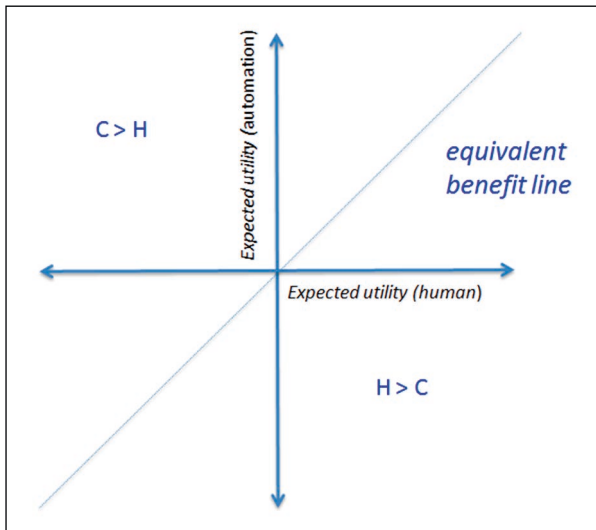


Fig. 4 A task space can be defined by the expected utility (EU) of executing a given task by a computer (C) and by a human (H). The equivalent benefit line is where $EU(C) = EU(H)$. Above that line there is greater utility in using technology, and below that line the utility is greater if completed by humans.

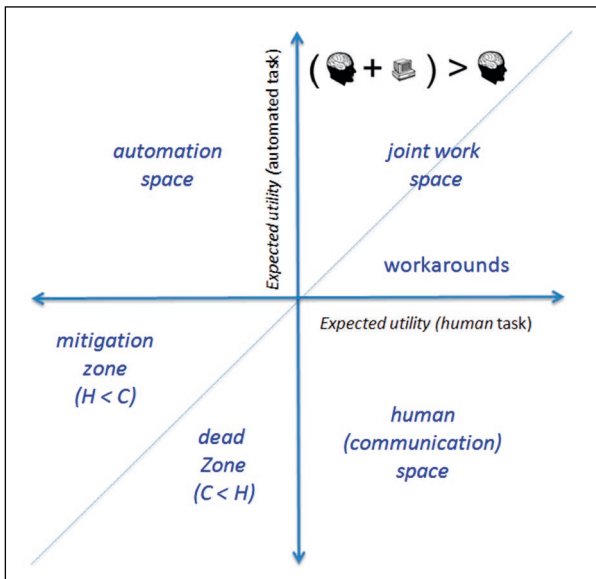


Fig. 5 The four different quadrants of the task space are better suited for automation, humans alone, for a partnership between the two, or best avoided.

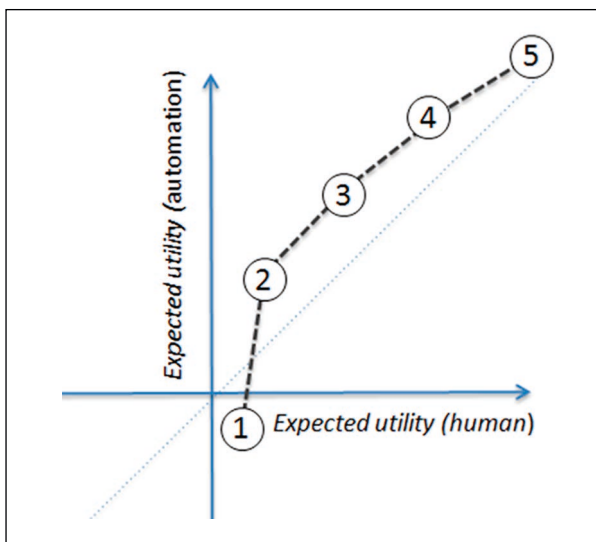


Fig. 6 For a given technology, the expected utility (EU) of completing a given task by human or computer can be plotted over task space. Figure 2 broke the information value chain down into 5 separate tasks (numbered one to five here). Here, the hypothetical profile for current generation electronic health records from Figure 3 is replotted into task space. The curve described by such plots is a function of the given task, the specific technology implementation, the human user, and the context of use. The shape of the plot varies by changing any of these four variables.

References

1. Ash JS, Berg M, Coiera E. Some unintended consequences of information technology in health care: the nature of patient care information system-related errors. *J Am Med Inform Assoc* 2004; 11:104-12.
2. Coiera E. Medical informatics. *BMJ* 1995;310 (6991):1381.
3. Friedman CP. A “fundamental theorem” of biomedical informatics. *J Am Med Inform Assoc* 2009;16(2):169-70.
4. Coiera E. *Guide to Health Informatics* (3rd Edition). 3rd ed. London: CRC Press, 2015.
5. Toth B, Gray J, Brice A. The number needed to read—a new measure of journal value. *Health Info Libr J* 2005;22(2):81-82.
6. Pluye P, Grad RM, Johnson-Lafleur J, et al. Number Needed to Benefit From Information (NNBI): Proposal From a Mixed Methods Research Study With Practicing Family Physicians. *Ann Fam Med* 2013;11(6):559-67.
7. Downs SM, Friedman CP, Marasigan F, Gartner G. A decision analytic method for scoring performance on computer-based patient simulations. *Proc AMIA Annu Fall Symp* 1997;667-71.
8. Fenwick E, Claxton K, Sculpher M. The value of implementation and the value of information: combined and uneven development. *Medl Decis Making* 2008;28(1):21-32.
9. Coiera E. When conversation is better than computation. *J Am Med Inform Assoc* 2000;7(3):277-86.
10. Dekker SW, Suparamaniam N. Divergent images of decision making in international disaster relief work. Ljungbyhed, Sweden: Lund University School of Aviation; 2005.
11. Coiera E. Technology, cognition and error. *BMJ Qual Saf* 2015;24(7):417-22.
12. Coiera E. Communication spaces. *J Am Med Inform Assoc* 2014;21(3):414-22.
13. Sintchenko V, Coiera EW. Which clinical decisions benefit from automation? A task complexity approach. *Int J Med Inform* 2003;70(2):309-16.
14. Tolkien JRR. *The Lord of the Rings: One Volume*: Houghton Mifflin Harcourt; 2012.

Correspondence to:

Enrico Coiera
 Australian Institute of Health Innovation
 Macquarie University
 Australia
 E-mail: enrico.coiera@mq.edu.au