COMMENTARY

The forgetting health system

Enrico Coiera

Centre for Health Informatics, Australian Institute of Health Innovation, Macquarie University, Sydney, Australia

Correspondence
Enrico Coiera, Centre for Health Informatics, Australian Institute of Health Innovation, Macquarie University, Sydney, Australia.
Email: enrico.coiera@mq.edu.au

Funding information
National Health and Medical Research Council, Grant/Award Number: APP10541465.

Abstract

Introduction: Forgetting shapes learning in two different ways. It impedes learning when important lessons are forgotten. Equally, it can be difficult to enact new lessons if we do not let go of old beliefs and practices that are no longer useful. A learning health system (LHS) that wishes to improve health service delivery will need to find ways to remember processes that shape quality and safety - using data that often resides beyond electronic health records. An LHS will also need to "forget", or programmatically decommission, obsolete practices, whose persistence otherwise leads to unnecessary system complexity and inertia to change.

Discussion: New forms of data needed to improve health services include process metrics extracted from digital systems; human-level metrics that capture workflow patterns and clinician behaviors; and multivariate process patterns that can identify service "syndromes." To avoid inertia to change, system complexity must be reduced by retiring (or forgetting) inefficient or unhelpful work practices. Biological models of programmed cell death provide a rich set of mechanisms to decommission elements of health services. These models suggest health service elements should be able to detect the end of their useful life and should contain internal mechanisms to orchestrate decommissioning—in contrast to current service decommissioning, which is an externally initiated, top-down down-driven process.

Conclusions: An LHS should take advantage of digital infrastructure to bring together people, sensors, analytics, and quasi-autonomous mechanisms for service adaptation. By drawing inspiration from biology, we can design LHSs that do not just remember but also actively forget.

KEYWORDS

clinical inertia, complexity, decommissioning, apoptosis, standards

1 | INTRODUCTION

We should learn from the past. To not do so is to waste opportunity and resource on failed strategies and practices. This is the fundamental proposition for building a learning health system (LHS)—to harness the treasure trove of clinical data stored in electronic health records (EHRs) so that every patient's experience adds to the knowledge base,[1-3] and every new patient's care is as effective as we can make it.[4]

Beyond questions of diagnosis and treatment, an LHS should focus on improving the quality, safety, and effectiveness of health care processes.[5] There is little point in discovering the best way to treat a patient, if poor execution leads to unnecessary harm, cost, or suffering.

Yet today health care services everywhere still struggle with quality and safety challenges, despite a decade of intense focus on the problem.[6] There is too much variation in patient care, and too much waste and harm in the system.

2 | QUESTIONS OF INTEREST

"Forgetting" is the complement of learning; and in this paper, I explore two different roles that forgetting can take in the LHS. First, much of the data and knowledge that is needed to address questions of quality and safety is today forgotten, and will likely never be captured just in

clinical records. Finding ways to capture these process data will likely require a broadening of the LHS vision beyond the EHR.

Secondly, forgetting is not always negative. It is just as important to engineer ways of forgetting obsolete practices in an LHS, as it is to discover new ones. Coming to grips with these two aspects of a forgetting health system are likely to be major challenges in bringing forth the truly adaptive, self-optimizing health care systems of the future.

3 | DISCUSSION

3.1 | Process information is lost

Quality and safety are ultimately process-centric properties of a system. They are shaped both by system defences that constrain unsafe actions, and system affordances that give latitude to actions that may come with risk. If we are to build learning systems that minimise patient risk, then it will be necessary to measure process execution. The state of the art in quality and safety measurement, however, significantly lags our ability to measure patient outcomes. Critical incident reports, for example, capture only a small proportion of adverse outcomes, are not representative of true event frequencies, and are not available in real time.[7] There are emerging digital technologies that trigger alerts when certain high-risk actions occur, such as repeat orders for a medication within a short time window.[8] In general, however, we still do not necessarily know which processes of service delivery should be instrumented nor which events are the most important to flag.

As a result, much process data still sit in the heads of those who work within an organization.[9] The nuances of what was done, when it was done, and why the execution of events was in one sequence and not another stay on the shop floor. There is a significant difference between work as imagined (for example, in a documented treatment plan) and work as done (the real-time execution of that plan).[10] Responding to the evolving logic of events in the physical world, clinicians must find ways of satisfying multiple competing demands and will not always be able to reconcile what is recommended practice with what could or should be done.[11]

Creating a learning health service that can optimise process execution would require concerted effort to "instrument the enterprise" and capture service information at a number of levels:

- Automated process-level metrics: Treatment and diagnostic events recorded in the EHR are one source of process information. For example, time stamps on events such as the creation of a medication or test order, and the steps that follow as the order is executed, can provide valuable information on process quality and efficiency.[12,13] Process mining of human-computer interaction logs can provide rich information about the effectiveness of workflows and software systems. Nonclinical systems are also a significant source of process data. Telephone metadata including Global Positioning System positioning data and online social network engagements can be used to monitor population-level health services, emergency services, and to generate consumer-derived data—a rich source of quality and safety data.[14]

- Human-level metrics: Measuring what patients and clinicians do can be very revealing but currently requires significant investment in data capture, for example, through direct observation or analysis of video and audio records. New classes of sensor, such as wearable cameras, or contact and location sensors can make workflow patterns in health services visible and open to detailed analysis. For example, wearable proximity sensors allow hospitals to trace likely infection transmission routes in hospital wards as patients, visitors, and staff move about and interact, suggesting modified infection control procedures.[15]

- Process patterns: Making sense of the causation behind system performance is a major challenge, as many service problems have complex origin, and will not yield to simple statistical analysis. Individual metrics can only tell us so much, and learning how different variables interrelate is essential to modifying the behavior of what is a genuinely complex system. For example, "syndromes" or telltale multivariate patterns may signify the likely source of hospital information systems problems.[16] Process mining and other machine-learning methods underpin the discovery of such patterns.[17,18]

3.2 | Learning in complex systems also requires forgetting

In patient safety, there has been a dawning recognition that creating safe systems cannot rely simply on focusing on what goes wrong (Safety 1) but that there also needs to be attention to what goes right (Safety 2).[10] We probably need to see a similar change in framing for LHSs, moving from a focus solely on learning what is right (LHS 1) to also forgetting what was wrong (LHS 2) (Box 1). Just because the evidence shows a practice is obsolete does not mean it instantaneously disappears There needs to be a concerted effort to communicate the need for change, and then embed the change into work practices.

Box 1: A typology of learning health systems

LHS 1: A system with explicit systemic learning mechanisms characterized by the use of information to generalize lessons within the system

LHS 2: A system with explicit systemic learning and decommissioning mechanisms characterized by the use of information to both generalize lessons from within the system and maintain efficient system function through controlled decommissioning or forgetting

---

Terms such as decommissioning, disinvestment, reconfiguring, rationing, and de-adoption describe the process of planned removal of old, unwanted practices such as a type of surgery, investigation, or

therapy. Collectively, these processes try to identify "low-value care"—practices that are not supported by economic measures of value.[19-21] Decommissioning is usually a top-down intervention, involving change-management strategies such as community engagement, and the use of champions.

Standardisation also seeks to eliminate "unwarranted" local variation.[22] Whilst standardisation clearly has an important role, it poses a challenge to the LHS, because its intent is to suppress local adaptation, yet creating variation is the whole point of the LHS.[23] Another well-known challenge to standardisation is that not everything is a good target for homogenisation. It is no doubt a good thing for computer systems to share a standard way to describe data or to construct messages. It is less clear that the design of user interfaces should be identical across different clinical services, which have different workflows, tasks, and goals. The overemphasis on top-down standardisation at the expense of local variation has been a poor policy choice in the realm of health information technology.[24] So it is too with clinical practices. Procedural variation in surgery seems to confer versatility that allows surgeons to vary their approach, depending on the specific needs of a patient.[25] Centres of excellence can have very good reason to do things differently to other health services, because of the different mix of patients they see, the unique resources they have, and the accumulated local evidence supporting their specialised approaches.

Decommissioning and standardisation can have limited impact in the real world. We may know only too well what does not work, but all our attempts to avert history repeat may have little impact. Indeed, health systems display a remarkable reluctance to shift performance in response to imposed change. Several decades of effort to improve the quality and safety of health care, for example, have in most cases had only marginal success in reducing rates of harm and adverse events.[6]

This system inertia—the resistance of a system to change despite clear evidence that change is essential—is an emergent property of the structure of health services, and is likely a function of system complexity.[26] Put simply, the more dependencies there are in a system, the harder it is in general to change behavior. Further, complexity grows over time, as we accrete new practices but do not entirely abandoning the old. One solution to growth in complexity and inertia is to actively reduce system complexity, freeing up the system to flex and adapt. Discovering mechanisms to overcome system inertia through complexity reduction thus becomes a foundational challenge in the construction of any LHS.

Local variations thus have the tendency to accrete over time and add to system complexity. They can persist in the processes, protocols and built structures of an organization, and in the workarounds, customizations, and annotations that happen to physical spaces.[27] Important lessons are thus embedded in the physical structure of the organization, and the physics of the way people act within that structure. With time, the canvas of a new organization is overlaid with accreted experience, lessons learned, and adaptations directly embedded into workflow. These structural memories are not inert, passive, or idly awaiting analysis. Rather, they sit there every moment—shaping work, constraining behavior, and altering human perceptions, actions, and intent. The task of forgetting old practices is thus non-trivial, as many are never documented or described, but simply become part of the fabric of work.

3.3 | Programmed cell death as a model for the forgetting health system

How one approaches mindful forgetting in health systems remains little explored. Whilst we have blunt top-down strategies like standardisation or decommissioning, there are no obvious complementary bottom-up processes that remove unwanted local variations whilst preserving what is important. Equally, there is much still to be learned about the best way to implement these mechanisms so that no harm is done in the process.

Biology may be able to help, as it has provided organizational science many metaphors and insights over the years, some more powerful than others.[28,29] Biological processes can also provide us with a set of mechanisms that parallel the organizational challenge of forgetting the unwanted and simplifying the complex. Specifically, programmed cell death (PCD) has exactly these roles in the organism.[26] It targets cells that require removal because they are no longer functioning well—for example, in the removal of precancerous cells. The PCD is also crucial in homeostasis. In embryogenesis or organism development, PCD helps craft organ structure by shaping which cells should continue to grow and which must die—for example, creating the spaces between fingers. The biology of cell death is complex and includes at least three different mechanisms of apoptosis,[30] necrosis, and autophagy.[31,32]

From the point of view of organizational development and function, the specific molecular mechanisms of cell death are probably of less interest than the functional design of these different death pathways. What is of interest is that PCD is adaptive to circumstances, can work from bottom up to top down, and has evolved sophisticated machinery to minimise unnecessary harm to healthy parts of the organism—all very desirable properties for the LHS.

At a high level of abstraction, the machinery of PCD has the following general features (Figure 1):[31-34]

- There is a separation between signalling functions, which convey messages and execution functions, which terminate cells based on state information.
- There are different roles for signalling. Some signals trigger cell death, others signal permission to continue operation.
- Signals can be generated at any of 3 levels:
  - Macro: Top-down signals originate far from a cell and come from high-level control mechanisms (called the extrinsic pathway for apoptosis in biology). If a cell loses contact with its surrounding cells and environment, it self-terminates;
  - Micro: When a cell becomes dysfunction, for example through irreparable internal damage, it releases local signals in a bottom-up fashion, to trigger self-destruction (known as the intrinsic apoptosis pathway);
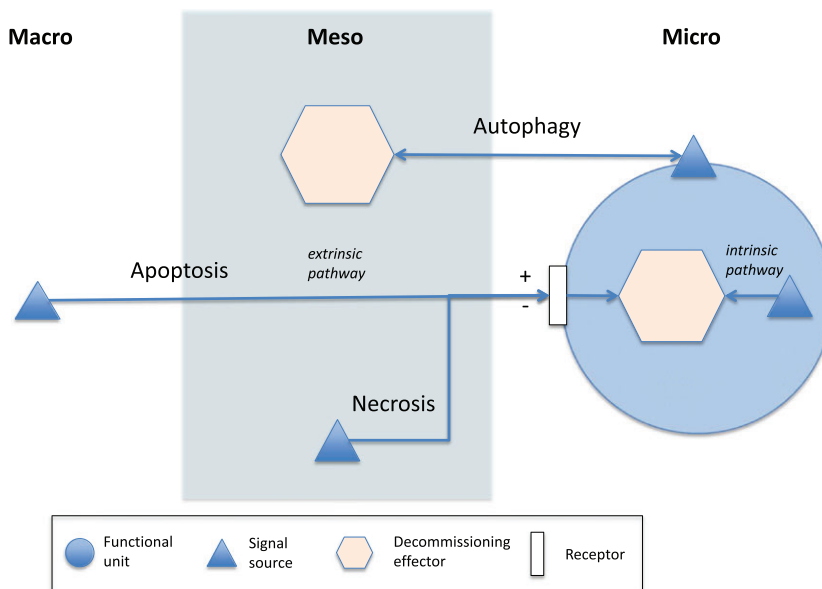
Macro   **Meso**   **Micro**

FIGURE 1 There are different control and execution pathways for programmed cell death including apoptosis, necrosis, and autophagy. These can be mapped into separate signalling and decommissioning mechanisms to manage unneeded functional units within an organization (however such units are defined). Separate signalling and execution mechanisms may exist at the level of whole of organization (macro), local external environment (meso), or internally to a unit (micro). Some signalling mechanisms share a common decommissioning process but may be external to a unit (extrinsic path) or within it (intrinsic path). Signals can tell a unit to continue functioning (+) or to decommission (−). Decommissioning machinery can sit outside a unit or within it.

○ Meso: "Middle-out" mechanisms include necrosis, which is triggered by events in the local external cellular environment; and autophagy, which is triggered by events on a cell's surface but co-opts distant cells to conduct cell termination;

• The different roles and sources of signals mean that a variety of signal receptor types and locations are needed. Cells contain both external surface "death receptors" to receive these external (macro or meso) signals, as well as a separate internal (micro) detection process;

• There is modularity in PCD design. Once a death signal is received by a cell from whatever source, a common internal mechanism executes the signal to die;

• There is redundancy in PCD design. Additional machinery exists external to the cell that can also terminate it. For example, if signals on its surface indicate that it appears to be dysfunctional, other specialized cells can destroy it (autophagy).

• There is variety in PCD design. Different classes of event trigger different pathways and mechanisms, and each class of cells may be regulated by a different bundle of pathways and mechanisms.

3.4 | Programmatic organizational decommissioning

What might biology teach us about programmatic forgetting in the LHS? The first observation is that there is a substantial gap between the elegance and richness in purpose and design of what we find in biology and current organizational mechanism. Where biology sculpts, organizations amputate and graft. The current health system approach to organizational forgetting is only top down. Decommissioning proceeds by telling us what is bad, and standardisation does it by telling us what is good. There is no dialogue between the local and top, as there is in biology, and there is none of the variety of mechanism nor local autonomy.

If we take biology as our guide, then for an LHS to be adaptive, the capacity to learn and change practice must happen at multiple levels from the local to the global. Whilst there is no obvious organizational equivalent of a biological cell, we can still talk about an organizational "unit," which operationally is substitutable with another similar "chunk" or element. Examples of units include a clinical guideline or a workflow; different units come together to constitute larger organizational "organs" like a hospital ward.

For programmatic organizational decommissioning (POD), we would require that each organizational unit be designed with its demise in mind. Each unit must be able to determine whether it should continue to operate or should terminate. It should also ideally contain the machinery for that termination. There should be clear mechanisms for the local to signal the central, and vice versa, so that there can be an ongoing and emergent "discussion" about where change is needed, and what needs to be retained. Box 2 contains some examples of the trigger rules one could build for different units, or aggregations of units, using the machinery of POD at different levels, from bottom up to top down.

In biology, rogue cells that fail to terminate can become cancerous. To minimise this risk, there is redundancy in design of PCD, so that if one mechanism fails, there is a good chance such cells will be caught by an alternate mechanism. In organizations, rogue units (such as a particular clinical practice) might also incorrectly persist and proliferate. POD thus requires backup mechanisms, just as with PCD, to police for such dysfunction. It should be possible within a digital infrastructure to monitor the process data generated by functional units to determine whether or not they are performing well. We should be checking to see whether they respond to decommission signals, and if they do not, remove them using an externally imposed mechanism. We can imagine software agents combing such a network, behaving like "cyber-immune" cells, checking digital entities for credentials of good health, as well as the digital footprints of physical entities. Failure to prove good health could trigger a central response.[35]

> **Box 2: Programmatic organizational decommissioning**
>
> Simple examples of rules designed to trigger decommissioning mechanisms within or external to functional units:
>
> - Micro: When new elements are added to data entry forms or computer screens, other elements will need to be retired to avoid increasing complexity and reducing functionality, eg, (Delete/Hide/Archive/Deprioritize) me if my (error/usage/incompleteness) rate puts me in the bottom 1% of elements.
> - Meso: The arrival of a new clinical guideline should require old ones to retire, eg, Retire me if (a more recent guideline exists/recent evidence contradicts my content).
> - Macro: Legacy information systems accrete with time and can constrain innovation by limiting the choice of new systems and absorbing resources through maintenance costs, eg, Decommission me if my (cost of maintenance and usage benefit) is worse than that of a newer replacement.

4 | CONCLUSION

It has been said that "the challenge of the LHS may require a novel emergent science of large-scale learning systems best seen as an evolution from the science of information systems, through a science of cyber-physical systems, and ultimately to a science of cyber-physical–social ecosystems."[3]

Biological organisms have evolved deep interconnected systems for cellular signalling and action, which support the growth of the organism and differentiation into functional organs, and which deliver homeostatic balance in response to external changes. Our challenge with health systems does not stop with becoming better at remembering the past. Our larger goal should be to develop an LHS that takes advantage of digital infrastructure to bring together people, sensors, analytics, and quasi-autonomous mechanisms for service adaptation.

For a complex adaptive system, there is no learning without complementary forgetting. We thus need to move our conception of the LHS from one focussed just on learning, to one that also is expert at forgetting, from LHS 1 to LHS 2. Hopefully, by drawing inspiration from biology, we can design these socio-technical machines "to understand the process of design from within the system, to design a system that more or less designs itself"[36]

CONFLICTS OF INTEREST

The author declares that there are no conflicts of interest associated with the material contained in this article.

REFERENCES

1. Etheredge LM. A rapid-learning health system. *Health Aff*. 2007;26(2): w107-ww18.
2. Gallego B, Walter SR, Day RO, et al. Bringing cohort studies to the bedside: framework for a 'green button' to support clinical decision-making. *J Comp Effect Res*. 2015;4(3):1-7.
3. Friedman C, Rubin J, Brown J, et al. Toward a science of learning systems: a research agenda for the high-functioning learning health system. *J Am Med Inform Assoc*. 2015;22(1):43-50.
4. Friedman CP, Wong AK, Blumenthal D. Achieving a nationwide learning health system. *Sci Transl Med*. 2010;2(57):57cm29-57cm29.
5. Corrigan JM. Improving Quality and Safety. In: Grossman C, Powers B, McGinnis JM, eds. *Digital Infrastructure for the Learning Health System: The Foundation for Continuous Improvement in Health and Health Care: Workshop Series Summary*. Washington, D.C.: National Academies Press; 2011:81-85.
6. Braithwaite J, Coiera E. Beyond patient safety Flatland. *J R Soc Med*. 2010;103(6):219-225. doi: 10.1258/jrsm.2010.100032
7. Macrae C. The problem with incident reporting. *BMJ Qual Saf*. 2016;25 (2):71-75.
8. Jha AK, Laguette J, Seger A, et al. Can surveillance systems identify and avert adverse drug events? A prospective evaluation of a commercial application. *J Am Med Inform Assoc*. 2008;15(5):647-653.
9. Coiera E. When conversation is better than computation. *J Am Med Inform Assoc*. 2000;7(3):277-286.
10. Hollnagel E. *Safety-I and Safety–II: The Past and Future of Safety Management*. Burlington, VT: Ashgate Publishing, Ltd; 2014.
11. Jaén CR, Stange KC, Nutting PA. Competing demands of primary care: a model for the delivery of clinical preventive services. *J Fam Pract*. 1994;38(2):166-171.
12. Ong M, Magrabi F, Jones G, et al. Last orders—follow-up of tests ordered on the day of hospital discharge. *Arch Intern Med* 2012;172 (17):1347–1349.
13. Westbrook J, Georgiou A, Dimos A, et al. Computerised pathology test order-entry reduces laboratory turnaround times and influences tests ordered by hospital clinicians: a controlled before and after study. *J Clin Pathol*. 2006;59(5):533-536.
14. Coiera E. Social networks, social media, and social diseases. *BMJ*. 2013;346:f3007. doi: 10.1136/bmj.f3007
15. Vanhems P, Barrat A, Cattuto C, et al. Estimating potential infection transmission routes in hospital wards using wearable proximity sensors. *PLoS One*. 2013;8(9): e73970.
16. Ong MS, Magrabi F, Coiera E. Syndromic surveillance for health information system failures: a feasibility study. *J Am Med Inform Assoc*. 2012;20(3):506-512. doi: 10.1136/amiajnl-2012-001144
17. Mans R, Schonenberg H, Leonardi G, et al. Process mining techniques: an application to stroke care. *Stud Health Technol Inform*. 2008;136:573.
18. Mans RS, Schonenberg MH, Song M, van der Aalst WMP, Bakker PJM. Application of process mining in healthcare - a case study in a Dutch hospital. In: Fred A, Filipe J, Gamboa H, eds. *Biomedical Engineering Systems and Technologies*. Berlin: Springer; 2008:425-438.
19. Niven DJ, Mrklas KJ, Holodinsky JK, et al. Towards understanding the de-adoption of low-value clinical practices: a scoping review. *BMC Med*. 2015;13(1):1.
20. Garner S, Littlejohns P. Disinvestment from low value clinical interventions: NICEly done? *BMJ-Br Med J*. 2011;343(7):d4519.
21. Daniels T, Williams I, Robinson S, et al. Tackling disinvestment in health care services: the views of resource allocators in the English NHS. *J Health Organ Manag*. 2013;27(6):762-780.
22. Wennberg JE. Unwarranted variations in healthcare delivery: implications for academic medical centres. *Br Med J*. 2002;325(7370):961.
23. Coiera E. Stasis and adaptation. *Stud Health Technol Inform*. 2013;194:11-19.

24. Coiera E. Building a national health IT system from the middle out. *J Am Med Inform Assoc*. 2009;16(3):271-273.

25. Apramian T, Cristancho S, Watling C, et al. "They have to adapt to learn": surgeons' perspectives on the role of procedural variation in surgical education. *J Surg Educ*. 2016;73(2):339-347. doi: 10.1016/j.jsurg.2015.10.016

26. Coiera E. Why system inertia makes health reform so hard. *Br Med J*. 2011;343:27-29. doi: 10.1136/bmj.d3693

27. Coiera E. Communication spaces. *J Am Med Inform Assoc*. 2014;21(3):414-422. doi: 10.1136/amiajnl-2012-001520

28. Hannan MT, Freeman J. The population ecology of organizations. *Am J Sociol*. 1977;929-64.

29. Penrose ET. Biological analogies in the theory of the firm. *Am Econ Rev*. 1952;42(5):804-819.

30. Lockshin RA, Zakeri Z. Programmed cell death and apoptosis: origins of the theory. *Nat Rev Mol Cell Biol*. 2001;2(7):545-550.

31. Bialik S, Zalckvar E, Ber Y, et al. Systems biology analysis of programmed cell death. *Trends Biochem Sci*. 2010;35(10):556-564.

32. Fink SL, Cookson BT. Apoptosis, pyroptosis, and necrosis: mechanistic description of dead and dying eukaryotic cells. *Infect Immun*. 2005;73(4):1907-1916.

33. Ashkenazi A, Dixit VM. Death receptors: signaling and modulation. *Science*. 1998;281(5381):1305.

34. Lavrik IN. Systems biology of apoptosis signaling networks. *Curr Opin Biotechnol*. 2010;21(4):551-555.

35. Coiera E. Information epidemics, economics, and immunity on the internet. *BMJ*. 1998;317(7171):1469-1470.

36. Coiera E. Putting the technical back into socio-technical systems research. *Int J Med Inform*. 2007;76:S98-S103.