



HEALTH CARE IN THE INFORMATION SOCIETY

VOL. 2

**FROM ANARCHY OF TRANSITION TO
PROGRAMME FOR REFORM**

DAVID INGRAM



<https://www.openbookpublishers.com>

©2023 David Ingram



This work is licensed under an Attribution-NonCommercial 4.0 International (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work).

Attribution should include the following information:

David Ingram, *Health Care in the Information Society, Vol. 2: From Anarchy of Transition to Programme for Reform*. Cambridge, UK: Open Book Publishers, 2023,
<https://doi.org/10.11647/OBP.0384>

Copyright and permissions for the reuse of this image is provided in the caption and in the list of illustrations. Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

Further details about the CC BY-NC license are available at
<http://creativecommons.org/licenses/by-nc/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Digital material and resources associated with this volume are available at
<https://doi.org/10.11647/OBP.0384#resources>

ISBN Paperback: 978-1-80064-803-6

ISBN Hardback: 978-1-80064-805-0

ISBN Digital (PDF): 978-1-80064-806-7

ISBN Digital ebook (EPUB): 978-1-80511-191-7

ISBN XML: 978-1-80511-192-4

ISBN HTML: 978-1-80511-193-1

DOI: 10.11647/OBP.0384

Cover image: D koi. *White H₂O Molecules*, June 27, 2022,
<https://unsplash.com/photos/5nI9N2wNcBU>

Cover design: Jeevanjot Kaur Nagpal

8½. Halfway Houses towards openCare–Stories of GEHR, openEHR and OpenEyes

This half chapter introduces missions and movements that have evolved from adventure of ideas, through anarchy of transition, into central components of programmes for reform of health care services, now extending across the world, at scale. It is not a pitch for their adoption or a comparison with other endeavours. It is an eyewitness account of how they came to be, and a perspective that has unfolded alongside of what the future might be created to look like. It is these aspects that seem important to record, so that progress can continue to be made.

The principal story told is of a mission to help bring coherence to electronic care records. This is the story of GEHR and openEHR—persisting along a thirty-year stretch of my songline. Its survival and continuity have rested on the enduring commitment of its pioneers and a growing, vibrant, humanly variegated (and sometimes quarrelsome!) community of creative and determined participants. It has had stalwart friends and supporters but, until quite recently, enjoyed almost negligible public funding. It is an iterative and incremental story of implementation that has embraced new perspective, approach and delivery of digital care records. I have described the three top priorities of openEHR as implementation, implementation and implementation. Only by enacting such vision can one learn how to do it. As Robert Oppenheimer wrote in his immediate post-war Reith Lectures, which I referenced in the book's Introduction, in attempting such a mission we discover who we are. The second story, told in less detail and combining with the profile of its founding pioneer, Bill Aylward, in Chapter Eight, is of OpenEyes. This initiative has evolved and disseminated a state-of-the-art open-source eye care record, now supporting around fifty percent of ophthalmology services in the UK. It has been made possible by a public sector-led collaboration of clinicians, NHS Trusts and companies.

Care records are concerned with capturing the 'Who did what, when, where, how and why?' in support of the health care of individual citizens. This half chapter seeks to encompass these same attributes. It is a story of the creation of halfway houses that have been instantiated today, along

a path creating common ground on which the future care information utility can grow in the coming decades. The mission to imagine, create and sustain this coherent, citizen-centred, well-governed and trusted resource will be central to future health care, as the world turns upside down in transition from Industrial Age to Information Society.

If trillion-dollar funding streams had been utilized differently, the kinds of mission described here might have saved the world much money, heartache and lost opportunity. Enacted faithfully and well, positioned at the centre of the care information utility that they can now help to create and sustain, such missions will contribute shared common ground that enables the world of health care to become a more caring, equitable and sustainable place.

This book has attempted a novel history, its structure inspired by Julian Barnes's novel, *The History of the World in 10½ Chapters*, hence its ten chapters and this half chapter.¹ Each chapter brings new and complementary account and perspective. The whole might well be described as a history in ten-and-a-half books! They are woven together along both historical timeline and personal songline.

Barnes's book starts with a stowaway woodworm's story of Noah's Ark and builds further stories in successive chapters, each based on events and interpretations that cast a different light on the history of the voyage, and therein on the ambiguity of all histories. It is a very entertaining mix of complementary stories—some fantasy, some historically sourced with precision, presented and interpreted with a delicious mix of intelligent and insightful commentary, understated ridicule and amusement. I knew Julian as an undergraduate at Magdalen College, University of Oxford, and was thus drawn to and loved reading his books. His writing feels perfectly in tune with the person I met there; a penetrating thread of acerbic discord runs throughout. He probably likes people like that—I think I recognized a clinician of similarly dynamic and penetrating personality, who I also knew a while ago, acknowledged in a cameo role in the book.

Barnes entitled his half chapter 'Parenthesis'—it is a personal and touching story of the ambiguity of love. That is how I read it—he did not say it as such. My half chapter, also numbered Eight and a Half, is a personal story of adventure, anarchy and reform, played out in two movements building towards common ground on which to base a future care information utility. These are the openEHR care record platform and its associated clinical information models, and the OpenEyes clinical ophthalmology electronic

1 J. Barnes, *A History of the World in 10½ Chapters* (New York: Knopf, 1989).

medical record application. They are building on the legacies of other pioneering initiatives that the book has connected with and drawn together. Legacy is fundamental to reform.

I have framed my half chapter like this, in part to emphasize that it is my story, and that others will, rightly and appropriately, have theirs, as players also closely involved in or observant of the events it describes. Stories of eyewitnesses bring historical perspective to the playing out of anarchies of transition and the unfolding of programmes for reform. How did the story I relate come about and unfold? The world will give different answers. This account represents my own experience, as someone who has been involved, from the beginning, with varying degrees of direct involvement: capturing ideas and designs; building and leading core teams and implementations; creating and sustaining interdisciplinary and multiprofessional environments, partnerships and alliances needed for the work; navigating storms and resolving conflicts; and establishing and managing the legal frameworks required for wider dissemination, governance and growth of the missions described, within worldwide context.

In terms of the credit and intellectual property embodied in these now very wide-ranging endeavours, they truly are the work and accomplishment of all the brave and committed souls who have participated in creating and sustaining them. They are heroes. Those closest to me feature and are acknowledged widely throughout the book and in its archive of additional resources.² Many hundreds of others now populate the websites of openEHR and OpenEyes.³ Perhaps the most important aspect of each generation's successes is demonstrated in the strength and staying power of the succeeding generation of heroes that it, in turn, enables and inspires. This half chapter is, in part, also an acknowledgement that the missions it describes are, as yet, still halfway houses along the road to an information utility that supports the reinvented and recreated health care services of the future.

Legacy and Reform

The term legacy has negative connotations in the context of information technology (IT), where it is often associated with incompatible and out-of-date systems that impede progress. Let us be more positive, here. Legacy is what we create and re-form from what we inherit, and then pass on, to be built on with new ideas and in context of new requirements—what

2 Available at <https://www.openbookpublishers.com/books/10.11647/obp.0335#resources>

3 *openEHR*, <https://openehr.org/>; *OpenEyes*, <https://openeyes.apperta.org/>

speaks for us when we do not, or no longer speak. A good legacy shows a way forward and helps create community and environment to carry the continuing burden of reform. Where there is a good way to follow, the will to follow it will grow. A bad legacy gets in the way. In the Information Age, we have become aware of legacy information systems that sink costs, fail to or no longer work well enough, and block future reform.

Legacy and programme for reform appear and connect in many guises. Their histories embody ambiguity, just like Barnes's stories of Noah's Ark. They are inherited as preface and passed on as postscript. As we play with words, they evolve, enriched by new conceptions, descriptions and inscription of ideas, informed by iterative and incremental experimentation and the experience we gain. They extend and connect over time: through prescription, by way of future standards of theory and practice; through subscription, by way of resources invested and governance applied; and through trust, by way of nurturing of community. The formative stages of reform depend especially on trust, which must be grown and sustained. Trust connects and cascades throughout programmes for reform.

Legacy also connects with law, by way of legates and legality—good and not so good people, good and not so good law. Reform connects with rebirth—reformation with renaissance. In the world of health care IT, our legacy is what we have helped to implement and sustain—good bits and not so good bits. Health care connects all around the circle of knowledge and throughout individual lifetimes of experience. It is in the accumulating common ground of this knowledge and experience that such wisdom as any may possess, resides, and holds its value and meaning.

Legacy and reform encircle past, present and future Dreamtime-like realities. They convolve with **songline**, **landmark**, **environment** and **trust**. These all feature, constructively and destructively, in the connections made. Those emboldened letters contrive **solvent**; legacy as solvent, ideas as solute and reform as solution, perhaps! Legacy as medium that enables and sustains people and ideas to grow and support future life; like water—the image of its molecules being the cover image of this book. Solute dissolves in solvent to make solution. Solutions arise in the mixing of solvents and solutes.

New solutions become new legacies, passed on and enduring, beyond personal songlines into new and different environments and to participants in these new worlds. Legacy encompasses not only traditional forms like money and property, but also information in the form of useful knowledge, method, histories and stories—stories of people and lives, the causes they have served and carried forward. Legacy is about luck as well as intention, both good and bad. Times of transition encounter many sliding doors, and lives are cascades of many transitions. A mixture of goals is tackled

with a mixture of motives, and with a blend of methods, resources and understanding. Legacy, like information, has multiple connections over time; it conditions and is conditioned by events.

Much of what we struggle to achieve and put in place dies quickly when we are no longer there to keep it alive. Not because it was wrong or misguided, although that may be the case, but often because it is no longer in tune with the times or has become obsolete or no longer relevant. How can we best make our endeavours connect and pass on a useful and sustainable legacy? We owe this to the future generations; we have been extraordinarily lucky to live through our age. What each of us, individually, contributes to legacy and reform is for us to justify to ourselves and for others to work out and decide on.

I have much enjoyed and appreciated my part in the stories of openEHR and OpenEyes, one in large part and one in much smaller part. I have been with them from their origins, planting and helping them to grow, making connections, staying the course within their growing communities of endeavour. This half chapter draws together the story of the seeds of openEHR, the seeds of OpenEyes to which they connected, the ground in which they germinated, sprouted and grew from flimsy saplings to sturdy trees, and how they are seeding, cross-fertilizing and expanding into forests, in what are now world-spanning movements. They are mutually complementary and reflect two co-dependent concerns—one of platform infrastructure and one of application utility. I hope and expect these will, before long, be connected in one much wider story.

Opportunity Knocks

In Chapter Four, I described three sliding doors through which I stepped after my twenty years in the Department of Medicine and then my own small Department, at St Bartholomew's Hospital (Bart's), from 1975 to 1995, where I had focused on the mathematical modelling of clinical physiology and its application to intensive care medicine and the creation of computing resources for medical and pharmacy education. In that chapter, I outlined the formative context of these three transitions, and what followed from them, first at Bart's and then at University College London (UCL). In this half chapter, I draw together the events that unfolded through the second sliding door, as I took on leadership of the European Union Good European Health Record (GEHR) Project and its pioneering focus on the standardization of care record architecture. This project and my previous work with Jane Dacre and Maggie Nicol, in creating the Bart's innovative joint medical and nursing Clinical Skills Centre (the first of the three sliding door transitions), led us

to the third sliding door, several years later. Through this door we moved together, as a combined health informatics, medical education and health services research group, to the Whittington Hospital campus of the UCL Medical School. There, the second half of my career in health informatics and multiprofessional education unfolded from 1995, as further described in Chapter Nine. It was at UCL that the mission of openEHR crystallized and came into being over the following decade.

In 1990, Sam Heard, my colleague at the Bart's Medical College, approached me to lead a new research-based consortium and bid for funds to develop a common generic architecture for electronic health care records, within the extensive EU Advanced Informatics in Medicine Initiative (AIM). Together with Alain Maskens, an oncologist who had left his clinical practice to establish a small company in Belgium and develop a health care record product called HealthOne, Sam was instrumental in drawing together a group of industry, health care and academic partners for this bid. I had recently become the first professor in the United Kingdom in the emerging field of medical informatics, and between us we established the research workplan and were successful in a highly competitive bidding process for funds to take the project forward. How was this architecture going to be imagined, created, developed and sustained? I describe here the historic legacy of the AIM GEHR Project and connect it with the evolution of openEHR. Additional detail is provided in Annexes hosted in the book's archive of additional resources. It is an ongoing story.⁴

What we started in 1991–94, in GEHR, led over the following decade to a succession of EU research and development projects and commercial implementations led by Sam and Thomas Beale at the newly created Ocean Informatics company in Australia. Successive phases of an evolving clinical and technical architecture for the digital health care record were piloted and reviewed. They came to be known as the openEHR specifications. The section of these that deals with the generic models of clinical data structures, from which individual care records are constructed, now

4 To supplement the material in this half chapter, further details of people and ideas involved in the creation of the GEHR architecture and subsequent evolution of the openEHR methodology and mission are included as Annexes lodged in the book's archive of additional resources (available at <https://www.openbookpublishers.com/books/10.11647/obp.0335#resources>), as well as on the openEHR website. I aim to include in the book's archive, an updating storyline that corrects errors and omissions and provides an opportunity for others, too, to give an ongoing account of newly formative contributions. Further detailed lists of contributors can also be tracked on the openEHR website, mirroring the wider pattern of acknowledgement, these days, in some publications from large-scale projects, such as the Large Hadron Collider experiments at CERN in Geneva, where sometimes hundreds of participants are every day essentially involved in the work.

known as openEHR archetype models, has been incorporated within both the Comité Européen de Normalisation (CEN) and International Standards Organization (ISO) 13606 standard for electronic health record communication. The specifications governing associated openEHR platform software implementations, on which to host clinical applications software are freely accessible, used, and now form the basis of care record system implementations, worldwide. Their scope extends over care records, medicines management, clinical decision support tools, patient reported outcome measures, clinical workflow management and regional public health systems, such as for infection control. In addition, generic software tools support the lifecycle of the associated clinical models.

The openEHR specifications are model-based and can be used directly to generate code for openEHR-based platforms, applications and services. They seek to contribute to and underpin a coherent common ground of patient-centred care records. One that facilitates the semantic interoperability of digital care records, such that their meaning and context, as specified and vouched for by frontline care professionals, can be well captured, shared and communicated within and between different communities and levels of health care services, and their diverse specialisms, native languages, geographies and jurisdictions. Moreover, one that also contributes to, and underpins, rigorous technical interoperability of openEHR-compliant software products, such that they function together reliably and sustainably, within and between different technology implementations and vendor products.

The public domain openEHR specifications and associated internationally governed corpus of openEHR archetypes are foundational to the re-formed architecture of digital care records proposed in this book, which envisions a shared common ground on which to base a future ecosystem of coherent care information. Crucially, this must enable and facilitate a sustainable and citizen-centred information utility, evolving to meet the changing requirements of individual health care in tomorrow's Information Society. What I have called 'Care Information Utility with you in charge' or openCare.⁵

Going back to 1991, Sam caused me quite a jolt when he came to see me soon after we had been awarded the GEHR project grant, to tell me that he and his family would be moving back to Australia! He promised to maintain his commitment to the project, including through regular extended visits

5 CIU with uic: the palindrome appealed to me! Something like working both ways, from citizen to professional and professional to citizen. I have imagined a logo for this, rather as I imagined the name openEHR (=open air). Now, perhaps we might talk of the care information utility as openCare!

to London to work within the team we created and led, there and across Europe. He was good to his word and our close and trusted friendship has grown and endured in regular contact and exchange for now over thirty years. The GEHR project team at Bart's included: Lesley Southgate, a close colleague of Sam in East London Primary Care and subsequently the Professor of General Practice at Bart's; Dipak Kalra, another East End general practitioner (GP), who was, thereafter, instrumental in connecting the evolving work with standards bodies in Europe and internationally; David Lloyd (1940–2023), an early pioneer of medical electronics and electrophysiology signal processing at Bart's; and Marcia Jacks, who started her career as a medical secretary at Bart's and became our team administrator. Introduced to us by Jane Dacre, Sian Griffiths, a trainee rheumatologist, took a sabbatical year from her professional training programme and joined the clinical team. Introduced to us by Jo Milan, Thomas Beale, an Australian engineer and computer scientist based in London, joined a year later, as did Stanley Shepherd, who had originally represented a GP system supplier, Update Computing, that was a member of the initial GEHR Consortium but then quickly withdrew from the project. Its elderly owner invited me for dinner at the London Savoy Hotel, to weigh me up, and decided not to commit the company to the GEHR mission.

Some years later, Sam and Thomas joined forces with others, including Peter Schloeffel (who I had met at conferences in Australia, when a visiting professor there), to establish the Ocean Informatics company (recently renamed as Ocean Health Systems). Ocean thereby formed a commercial test bed for the developing ideas, focused on implementation within the health IT industry and product marketplace. The company became an influential early bridge with many other companies, helping them to understand the ideas, see their way to adopting them within their own products and services, and become central players in the dissemination of what became known as the openEHR methodology.

The die was cast in those formative years, in those formative environments, by those formative teams—a die to press new perspective, approach and delivery of a shared common ground of clinically focused and rigorous conceptual coherence of digital care record structure, and to lay foundations for the future evolution of citizen-centred care information utility.

As described above and in Chapter Four, as the third of three sliding door moments in my career of the early 1990s, in 1995 I was invited by the UCL Provost, Derek Roberts, and Chair of its associated NHS Trust at the Whittington Hospital in North London, Helene Hayman, to move there, along with some nine of my Bart's team colleagues of that time, to establish and lead a new academic Centre that connected health informatics, medical

and multiprofessional health care education and health services research. The story of this new centre, christened CHIME (Centre for Health Informatics and Multiprofessional Education), is told in detail in Chapter Nine, as an example of the creation of an innovative new environment. The opportunity was created and orchestrated by the Vice-Dean of the medical school based at the Whittington hospital, David Patterson, and the Dean of the Medical School, John Pattison (1943–2020), advised by my long-term colleague at UCL, Mark Leaning, who was based there in the Clinical Operational Research Unit (CORU) Unit that I described in Chapter Four. Sam, Thomas and another colleague who later became central to the openEHR mission, Ian McNicoll, were appointed as honorary members of this academic department, until my retirement from academic life in 2011. As was Justin Whatling, who went on to key appointments in the Cerner and, more recently, Palantir companies. Don Detmer, Peter Singleton, Tony Shannon, Mark Leaning and Tim Benson were closely involved with us in those times.

The not-for-profit openEHR Foundation was established at UCL in 2003, to take the work forward in the public domain. The now very extensive related intellectual property is made freely available under Creative Commons license. The openEHR methodology has since then connected progressively with electronic health care record systems, products and services throughout the world, guided by iterative and incremental implementation experience. Its operations have now transferred to an independent Community Interest Company with directors elected from the communities of individual subscribing members, and industry and health care organizational partners. In all the steps towards establishing a sound legal footing for the openEHR Foundation, we had outstanding support from the UCL Business department and its chief executive, Cengiz Tarhan, his staff member, Renata Tarnowska, and our first employee as manager of the new openEHR International community interest company, Jill Riley. We also benefitted from substantially *pro-bono* support from major London law practices specializing in Intellectual Property law and Charity and Community Interest Company law (Reynolds Porter Chamberlain, Oliver Bray; Bates Wells & Braithwaite, Stephen Lloyd and Abbie Rumbold). When establishing the Community Interest Company, we had invaluable support from a Swedish colleague, Gunnar Klein, and the Apperta Foundation, chaired by my OpenEyes colleague, Bill Aylward, and its Chief Executive, Peter Coates.

Throughout these early endeavours, seeking to capture and articulate what became the openEHR mission, Sam and Dipak were principal articulators of clinical requirement, and Sam a key innovator as well, able to break new and emerging concepts through into working software

prototypes which could then be learned from, refined and extended by the technical experts who joined in, from their academic, health care and company bases in many countries. Thomas anchored the technical side of this evolution, to provide a rigorous and effectively coordinated ecosystem of specifications and tooling. Sam's colleagues at Ocean, notably Sebastian Garde, Chunlan Ma and Heather Leslie, made the massive contribution of the Clinical Knowledge Manager developed and used for the curation of the clinical models. There was wonderful and essential complementarity of approaches to, and synergistic endeavours in support of, the openEHR vision and movement—there was a core group, and now there are many more, of leadership roles and leaders joining efforts under this umbrella.

Both Sam and Thomas have put astonishing amounts of sustained personal commitment and life work into the mission. I characterized the qualities required as a combination of technical rigour, clinical engagement and trust. Of us three enduring founders of openEHR, Thomas has been the pillar supporting technical rigour, Sam that of clinical engagement, and I have sought to provide strategic vision and an inclusive home base and environment, to hold things together in a spirit of trust that has—at times precariously—prevailed through thick and thin! I consulted on and set out terms of reference for wider review boards to anchor the architecture technical specifications (ARB) and clinical models (CRB). The former took root and evolved well. The latter has had a trickier and more troubled pathway over the years—reflecting a world where frontline clinical collaboration and governance have proved harder to anchor successfully for the domain, than have the more technical aspects. Notwithstanding such reflection, those who stuck with and anchored all parts of the mission can be justly proud of what they have achieved. Even those who were hostile to, or unpersuaded by, the openEHR mission, have noticeably paid a good deal of attention to its detail and evolution. We could see details of the growing traffic to the hosted website as evidence of this! All told, not an easy mission to lead, especially when one's principal employment responsibilities were much wider and mainly focused elsewhere, nearer to home, as exemplified in Chapter Nine!

Germination of Mission

It is no accident that the pioneers of health information systems whose contributions I celebrated earlier in this chapter were at one in their presence at the coalface of health care professional practice, and all but one clinically qualified. This reflects that I had chosen to base my career in health informatics in a medical school clinical department and hospital

setting, as a non-clinical academic. There were opportunities at that time to work as an IT professional, developing and delivering a support service in a health care setting. These opportunities were, in the main, line managed not by clinicians but by hospital managers. There was also opportunity to work in an industry or consulting setting but there the focus was more narrowly on product development and overview at a distance from frontline care. Many followed an academic route in medical physics or computer science departments, again with narrower research focus on collaboration with clinicians working elsewhere. My choice was to work in the middle of a community comprising all the disciplines and professions that come together in academic medicine and hospital care and build from there. To follow this exploratory approach, my need was for alliance with colleague clinicians within the medical school, where my skills could align with their everyday concerns in education, research and practice.

From my early Bart's days, I set about building such personal alliances to help and support me in pursuit of my exploratory academic mission. This required much time spent and attention paid to listening and responding to a wide range of clinical colleagues, about their individual perspectives, interests and needs. Highly intelligent people and each with an angle on what informatics might have to offer them! I looked for win-wins, aligning my area of skill and interest with theirs. The scope for informatics innovation spanned medical education, clinical research and development, and more straightforward IT support for everyday work. No amount of words would help in connecting our missions in practical terms—the sole unifying paradigm was one of practical implementation. Taking on and making and doing things that were perceived as interesting and useful, and thereby learning more about them and how to use the computer in novel ways to explore and improve practice.

Many of my activities in the Bart's Department of Medicine, working with junior doctors intent on gaining their professional doctorates, had been of this kind, such as I have described in Chapter Four. Writing software to enact Huw Llewellyn's ideas about the diagnostic decision making process, based on mathematical set manipulations; creating a novel database to store and process Andrew Gorsuch's complex and extensive time-series data tracking genetic and immunological concomitants of diabetes; creating non-linear optimization methods to match the Mac Series respiratory model with intensive care unit (ICU) data, with Charles Hinds, and innovative problem-solving exercises for medical students and postgraduate trainees, using this model; extension and refinement of the Mac Series drug metabolism model for use in pharmacy education and research, with Giles Saunders at the London School of Pharmacy; development with Giles and Steve Jackson at Bart's, of a new simulation model of drug prescribing. Of course, few of

these proved of lasting significance, but as a means for me to experience and get to know these communities and worldviews, it was a good environment.

I reflected in Chapter Four on how this stage of my career came to a conclusion in 1989 and a new one opened up, in combining clinical skills training and informatics concerns in a new joint initiative and taking on leadership of the GEHR project. As Bridget Ogilvie, the wise head of the Wellcome Trust of that time, had advised me, when I told her about winning the EU GEHR project funding in 1991, I was faced with a choice between talking and writing about it and doing it. She smiled, saying that the breadth of ambition was so great that, if I took the 'make and do' route, no one with power to back it would wish or find themselves able to do so, but if we were successful, they would always, secretly, have been our friends! Rather typical of the reception of attempts to tackle 'wicked problems', in general, one might say! But it is pointless to be critical of such reality—not everyone is cut out to be a pioneer, and we would not get far if they were! We need educators, commentators and managers to keep pioneers honest, as well! More problems arise from those who imagine they can pioneer meaningfully in a particular health care setting, while living in a different and otherwise focused setting, and thus not sensitively connected with and experiencing the impact over time of the implementation of their ideas.

In the context of the new world that I moved into with the GEHR project, my general approach was much the same as in my earlier period at Bart's. It focused on iterative and incremental implementation as the basis of building and sustaining new alliances—in this case much more extensive and international ones—to gain understanding of the domain, to enable and guide the formulation of new approaches. My principal role became one of creating and leading an inclusive working environment, drawn widely from across complementary disciplines, professions, organizations and industries. This pattern followed on into the new centre we went on to create at UCL. It was a pattern where leaders from multiple complementary domains had a home base from which to lead a wide variety of initiatives, and new team members could develop to become leaders in their own fields of endeavour.

To emphasize the point about learning by doing, I recall commenting sometime later, to a group of visitors from the European Federation for Medical Informatics, that the three most important priorities of the openEHR Foundation were implementation, implementation, implementation—it stuck! Implementation comprising new approach and method, environment and team, and governance—a trifecta I explore in Chapter Nine. In such a historically perilous domain as the grand challenge of the electronic health record, it was, admittedly, a riskily adventurous approach, but it was the only way to engage constructively and significantly with this

implicitly interdisciplinary and multiprofessional domain. It could easily have heralded a downswing towards disaster, but, fortunately, openEHR survived along an upswing towards success. This path, however, was not without turbulence and personal cost and disappointment to some, along the way.

Before describing the adventure of ideas that started with GEHR and are now central to openEHR, I will nail some more personal colours to the mast. These have underpinned my vision of what the mission has been about, how I set it up, prioritize, and balanced the complementarity of contributions to its methodology, and led it during my time at the helm. It was from the outset an inclusive mission, growing from the coalface of clinical, technical and organizational requirements for care records, and embracing perspectives and participants from throughout health care communities, professions, academia and industry. This breadth of vision and community has served us well, even if sometimes portrayed as naively academic and esoteric! It now has considerable track record of impact, albeit with an inconsiderable record of words written about it in journals and books. It is sometimes criticised for that failing, being thus considered insufficiently academic. It was, thus, both too academic and not academic enough! That resonates with the way the world views efforts to resolve wicked problems, as rehearsed in Chapter Seven! One cannot please everyone!

It cannot all be wrong, though—I am somewhat in agreement with the Thomas Lincoln (1929–2016) school of thought in that regard (that more clinical data sometimes betokens less clinical effectiveness). But I do recognize that it might have been better for me to have said and written more along the way. From some perspectives, yes, but with perspective focused on creation of the nuts and bolts of a practically grounded and sustainable contribution to the grand challenge of the electronic health care record, I remain unsure. I couldn't do both, as Bridget Ogilvie had wisely advised. And given the challenge and uncertainty of the times, and the above-described imperatives of implementation, my adding to the noisy anarchy of the times might not have helped much. In the wider history of the health informatics field, flourished staking and advocacy of ideas, as yet unimplemented, has too often led in short order to a busted flush of cards then played. I was never much good at poker! Joining in might have hindered the unique opportunity I had been given, to anchor a vision, create and shelter a base camp and staging posts from which to pursue it, assemble and hold together a team of doughty mountaineering volunteers, and collaborate widely across the world with them and others who joined in the quest, in the incremental discovery and implementation of the openEHR mission. We had to learn how as we climbed, and it has proved an Everest-scale ascent! Climbing a difficult mountain has to be step by step, from

below; there is no way to think, talk or otherwise project oneself straight to the summit! And as introduced in Chapter One and further discussed in Chapter Nine, where it is described as 'leading from below', insider mediation is best led and conducted under the radar, and I often had that kind of role during my career.

During Covid restrictions, while shielding and since, I have put sometimes six hours a day and more, up to six days a week, for almost three years of my retirement, into this book. I hope this might be seen as some measure of atonement for having written too little before! I have never taken the step frequently taken by academic chiefs in my time, of placing their name on the publications arising from their teams. I suspect I might now be thought to be writing too much—such is the phenotype of advancing years!

GEHR—Perspective, Approach and Delivery

The GEHR team's first attempt to capture the scope and architecture of electronic health records was in 1991. The objectives were simply expressed on page one of the project brochure published through CEC DGXIII–C4 Health Telematics (AIM).⁶

A common electronic health care record architecture for Europe—evolving through cooperative development:

- the development should be based on an experimental methodology, starting from clinical needs and ethical and legal requirements
- the architecture should be in the public domain

A good first step in any such quest is to focus first on the requirements and these were expressed and reported in terms of:

- Requirements for clinical comprehensiveness, Deliverable 4 (144pages)
- Requirements for portability, Deliverable 5 (141 pages)
- Requirements for communication, Deliverable 6 (139 pages)
- Specification of functional requirements for clinical use, Deliverable 7 (85 pages)

⁶ I have digitized this brochure and included it as an additional resource of the book. My brief history of GEHR, written at the time, is also included in the Annexes to this half chapter posted there (available at <https://www.openbookpublishers.com/books/10.11647/obp.0335#resources>)

- Ethical and legal requirements of GEHR architecture and systems, Deliverable 8 (69 pages)
- Educational requirements of GEHR architecture and systems, Deliverable 9 (62 pages)

I have the slides for these presentations in front of me as I write, here, which express their essence succinctly. Here is the simple diagram I invented to crystallize the mission, which formed the basis of the expression of requirements for clinical comprehensiveness:

An overview of requirements for the digital care record

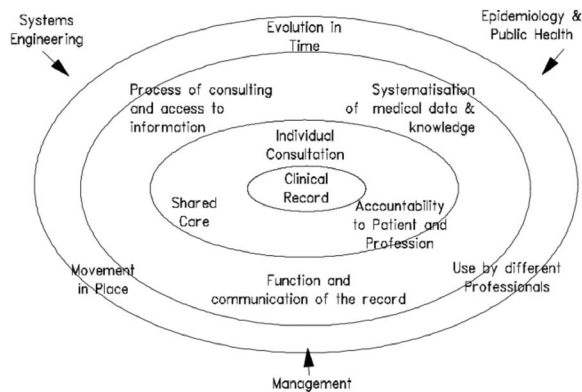


Fig. 8.22 The comprehensiveness of the electronic care record ecosystem, as expressed in the requirements for the GEHR architecture and mission. Image created by David Ingram (1992), CC BY-NC.



It focuses on ‘tripods’ or ‘trifectas’ of complementary concerns, in concentric ellipses, centred on the perspective of the citizen at the centre. This ‘thinking in triangles’ became embedded in my conceptualization and expression of the domain. In this perspective, a patient is present both as an autonomous individual and with their family and carers, in personal relationship with a clinical team, surrounded with information that is personal, shared with professionals and professionally accountable.

The next enveloping ellipse leads to formalization of information elicited and used within individual care episodes, combining with relevant knowledge of the domain, and set within an organizational ecosystem of health care services. And one step further outwards it connects with the evolution of care records kept, over time, in different places, and in interaction with different people and services. All this information is sustained and pieced together within the context of health system management, population health and technical infrastructure. The architecture we sought to characterize and specify was of the record of care produced and used within this whole ecosystem.

The diagram proved a good foundation on which to create, coordinate and lead the project. It was a quite simple framework that all could see and own—in the partner organizations of the project consortium and the Coordinating Partner project team that I built and led. It opened into a principled framing of the requirements that would underpin the architecture, ranging across the ecosystem described. This framework of requirements was taken up, almost entirely, in subsequent European and then International Standards Organization statements of requirements for electronic health care records of the era.

Here is what we drew together into the final description of the project, from which the following quotations are taken.⁷

7 *The Good European Health Record (GEHR) Project A2014, CEC DGXIII—C4 Health Telematics (AIM)* (Brussels: European Commission, 1994), pp. 1–16. My geometrically tutored mind often connects and visualizes issues in threes and triangles. Diagrams where I picture the tripod or trirecta of clinical, technical and organizational dimensions of coherence of clinical data populated many slides that I used to explain the GEHR project and successor openEHR missions, from the early 1990s. My initials being DI, these, and other geometrically inspired designs I frequently used, including one for challenging and inspiring my doctoral students to capture their contribution and its disciplinary context in a single image, became known by them as DI-agrams! In my mathematician days, I loved learning and taming geometry; for me it is a wonderful synthesis of the mathematical and the visual, and their connections with the natural and human world. There are other dimensionalities than threes, of course. Triangles tessellate, as do rectangles and hexagons. I started also to use hexagons to give visual impact to the network connectivity of the GEHR project's architecture of the digital care record. In Chapter Nine, I reflect on monisms as principles, dualisms and dichotomies as complementarities and choices, and tripods and trirectas as stable building blocks of the systems and processes required for imagining, implementing and sustaining a viable ('livable') organic and evolving care information utility.



Fig. 8.23 Front sheet of the GEHR Project Brochure, 1994. Image created by the GEHR project team (1994), CC BY-NC.

The GEHR project was established to develop a comprehensive and widely applicable common data structure (architecture) for using and sharing electronic health care records in Europe. The information environment in which such an architecture would be applied includes all sites capable of creating and maintaining medical and related data. It encompasses many different system types, networks, database types and vendors, and also many levels of software engineering capability. The user organizations range in size from large health regions and hospital groups with dedicated computing departments to single handed practitioners with a PC and a modem.

Clinicians are becoming increasingly aware of the opportunities the computer could offer to support their clinical practice. The management of complex diseases, clinical audit, and the automatic generation of reports are examples. The growing complexity of health care provision means that resource managers need greater access to aggregate information about the processes of clinical care. Unfortunately, the computer systems currently used in most hospitals and general practice surgeries, and more importantly the data modelling concepts which underlie them, are ill equipped to cope with these new challenges. The Good European Health Record architecture provides a framework which supports the full diversity of clinical data storage and communication required by clinicians. It is formulated to encompass the different disciplines of primary and secondary health care, for doctors, nurses,

and other professionals and in all European countries. Ready access to a wide range of datatypes is of increasing clinical importance, and the work of the project has included these multimedia aspects of the record architecture. Examples specifically addressed include X-ray and photographic images, bio-signals, technical drawings and most importantly textual information, for example clinical observations and laboratory data in the form of coded terms and free text.

The architecture has been derived on the basis of a full analysis of the requirements for an electronic health care record to support ethically and legally acceptable individual patient care. The GEHR project has employed iterative prototyping by clinicians to explore and test evolving ideas in practical situations. The project has consulted widely and has sought to establish and maintain close working relationships with other projects and teams. The project has members on CEN TC251/WG1/PT011 and the GEHR work, and deliverables have contributed the early standardization work being coordinated there.

On completion, GEHR will offer:

- architecture description
 - formal object model
 - exchange format
- multilingual dictionary of health record items
- library of anatomical drawings
- specifications for data access and integration tools
- a follow-on initiative to support future development

These results will be made available in the public domain.

GEHR is an architecture with supporting data sets, specifications, and recommendations for the implementation of compliant systems. GEHR is not a medical records software system.

Requirements for clinical comprehensiveness

The foundation of the development of the GEHR architecture was a thorough review of the clinical requirements for recording and processing patient information. This work involved literature reviews, questionnaire surveys, and group discussions, and was supported by the evaluation of successive prototypes. This preparatory phase covered the first year of the three-year project. The priorities of the GEHR project have reflected the belief that the clinical record is most necessary, and should be most available, when a clinician is offering care in a consultation. Thus, any compromise should always be directed towards offering quickly

accessible, accurate, and complete information to an authorized carer when attending a patient. The clinical record will be used by staff trained in different disciplines, working in different settings, on different sites and in different languages. The architecture must facilitate record storage on different sites and provide a common interchange format between heterogeneous systems. It must accommodate evolving needs for coding and classification standards and for the use of clinical guidelines in the management of care. The clinical record must accept these three areas of change: in time, place, and clinical perspective. A health care record evolves gradually over a person's lifetime, and family records over generations. A person's health care needs will change and evolve in time, as does the practice of medicine, and the economic and social framework within which medicine is practised.

Ethical and legal requirements

Ethical issues are fundamentally important because the use of electronic health care records (EHCRs) brings a risk of serious harm to patients or clinicians. However, the risk can be minimized without compromising the usefulness of the record, and regulation is both technically feasible and morally appropriate.

As understanding of many of the ethical issues depends on understanding of the purposes of the EHCR, these have been made explicit. These purposes have been assigned to a hierarchy which will itself aid the resolution between competing ethical imperatives. The primary purpose of the EHCR is to benefit the patient by providing a record of care which supports present and future care by the same or other clinicians. The secondary purpose is to provide a medico-legal record of the care provided and hence to demonstrate the level of competence of the clinicians involved. Tertiary purposes must be legitimate (involve consent) and can never be allowed to compromise the primary or secondary purpose. Examples of tertiary purposes are the generation of data for health service management or public health programmes.

Two important foundations of the relationship between a clinician and a patient are the delivery of clinical care to the highest standard and respect for patient autonomy. The latter inevitably leads to a proposal that the right to informed consent and the right to confidentiality are also moral principles of the highest importance underlying implementation of a 'good' EHCR. Patients should exercise as much choice over the content and movement of their medical records as is consistent with good clinical care and lack of serious harm to others. Records should be created, processed, and managed in ways that optimally guarantee the

confidentiality of their contents and the legitimate control by patients over them. The record must be secure yet accessible to patients.

The project has also considered the legal principles which have a bearing on the EHCR in terms of confidentiality, ownership and copyright, liability and accountability, identification, durability, processing of personal data, and transparency. The present diversity of legalisation on these issues is uncoordinated. There will be a need to harmonise legislation if movement of medical records is to be sanctioned by clinicians and patients.

Ethical and legal acceptability:

- preserve patient confidentiality
- respect patient autonomy
- faithfully record clinical actions
- only allow appropriate user access
- facilitate adequate audit trails and backup

Requirements for education

All health care students will need to be familiar with operating electronic health care records and with the ethico-legal framework in which they must be operated. It is therefore proposed that a portion of the electronic health care record is dedicated for student use. This would enable students to gain experience in making records, but this portion would be excluded from service functions. The record must also support teaching of students and aggregation of data for educational purposes.

Requirements for portability

Portability independent of:

- Hardware
- operating system
- software application
 - database, network, programming language
- national language
 - coding system

The GEHR architecture seeks to support clinical records which are independent of hardware, operating system, software application and the language used to record the clinical information. Language independence includes not only national language but also medical language and coding systems for medical language. Language independence clearly requires

the translation of the contents of clinical records. But there is widespread agreement that the electronic health care record must be structured and that the meaning of the data may depend on their context. GEHR clearly should not favour any one national language, nomenclature or coding system over another. The original language must be identified with the health care data in the EHCR itself.

Requirements for communication

The GEHR project has reviewed emerging clinical and technical standards for the communication of health care data. The project has developed a consistent approach to deal with the functionality and capacity issues encountered when incorporating 'bulky objects' and other externally held data into the health care record, covering the following communication requirements:

- standard representations for data types (images, ECGs)
- standard data sets (laboratory data, drug prescriptions, minimum data sets)
- Standard messages and data transmission protocols
 - standard EDI and ODA protocols
 - confidentiality of data transfer

Technical functional specification

Preservation of meaning

Clinicians value the facility for individual expression and creativity within the EHCR; however, this may make it more difficult to share. In order to ensure that meaning is preserved when the record is transferred from one computer to another, information should be recorded within its context. The original language and terms set should be identified, and original views of the data should be retained to maintain the grouping of specific pieces of information.

The boundary of the record

The clinical record must be clearly defined, and information should not form part of the record until a clinician has taken responsibility for that information and placed it into the record. Information within the computer system must be held in something like an electronic mailbox, and only considered part of the record when it has been committed to the record by an authorized person.

The transaction

In order that electronic health care record may grow logically in a way that preserves its integrity, and complies with ethical and legal requirements, it is proposed that the transaction should form the basic unit of the clinical record. In fact, the clinical record may be considered as a set of such transactions. Within the GEHR project, a transaction is defined as 'the information recorded about a patient by a single author in one institution at one point in time'.

Description of the GEHR architecture

The results from the foregoing requirements deliverables provided the basis of a first attempt to define a formal data architecture, in largely clinical terms: the Interim GEHR Architecture. This had the objective of providing sufficient flexibility to accommodate all the potential individual styles of record and define a set of constructs with which to model the data and concepts used in clinical practice. The fundamental architectural components evolved by GEHR for specifying what is contained in the record may be summarized as:

1. The *Transaction*
2. The *Health Record Item*
3. The *HRI Collection*

Each of these is further defined in terms of attributes which address aspects of identification, content, and context.

Every effort has been made to derive the most generic, flexible, and prescriptive structure possible. But where conditions have identified the need to be prescriptive (for example in situations where medico-legal security must be maintained) the model incorporates features which may be utilized for this purpose. The EHCR is the top-level containment structure and would be composed of many transactions, together with some data enabling the record to be identified.

A key specification of the clinicians within GEHR has been to treat the clinical encounter as a special grouping of data items for medical legal reasons. This grouping, termed transaction, has been fully documented in the functional specification. It reflects the data entered in one interactive session with a patient record—either a consultation or perhaps the 'filing' of a test result or letter. Common transaction identifiers might also form part of the context characteristic and allow a complete consultation to be identified and processed as one unit of the computer system. This would be used, for example, during the transfer of information to another institution.

The *Transaction* is a containment structure for collections and/or items which are committed to the record by an authorized person at a particular time and place. Transactions do not contain other transactions.

The *Health Record Item* has been proposed (and has been adopted by CEN PT011) as the basic unit of health information within the record. This represents the finest granularity by which an individual piece of information may remain meaningful if viewed in isolation (although complete interpretation may require it to be seen in perspective of other health related items—the clinical context). In essence, the Health Record Item is composed of an item name, its primary content value, and other associated identifiers, properties, and attributes. ‘Weight–76 Kg’ and ‘Family history–Hypertension’ are simple examples.

Health Record Item Collections allow for the construction of more complex aggregations of data. Examples are the decomposition of ‘blood pressure’ into ‘systolic’ and ‘diastolic’ components, or the breakdown of an antenatal examination into several sections. The recursive structure of the collection allows the health record items to be assembled into completely flexible but valid structures, of which the largest collection would be the entire patient record itself. The overall item and content values within it can each be further elaborated. The *Content properties* are used to further define the content value, such as units. The *Context characteristics* are features which relate to the whole item (its name, content, and properties) such as date and time of recording, author identification, language used in the recording.

The focus of work of the project moved to concentrate on developing from the Interim GEHR Architecture towards defining a comprehensive and rigorous information model appropriate to the content of a multi professional, multimedia health record. One particular application of this model is in deriving a formal view which can constitute an exchange format whereby safe and rigorous exchange of clinical data may be undertaken. The end objective is a formal model of the data defined in terms of object classes and structures which capture the full semantic richness of the clinical and ethical legal requirements. This is likely to be an important requirement for a formalism to anchor the future common health record architecture for Europe, with capability for monitoring of compliance to specified standards.

GEHR thus embraces an architecture which can be used to define the progressive adherence to standards for the clinical content of records and for compliance with requirements for ethical legal practice.

Object model and exchange format

The GEHR project has developed two formal definitions in support of its proposals for a common electronic health record architecture. The GEHR

object model defines the structure and content of information at a site. The GEHR exchange format defines the information exchange between sites.

It is recognized that what needs to be standardized is not how people practise medicine, but some useful minimum semantics of the information recorded in the process of care. Thus, the most basic aim of the information model is to enable efficient and effective computerization of existing and new medical information, not to suggest or prescribe how practitioners should do their work. At the technical level, the model of the information held at a site is most importantly a model of 'standard' underlying information structures and not a model of any particular view of such information as seen for example on a screen, on printouts, or as grouped within particular database methodologies. This is because in the heterogeneous European health context there are a multitude of users using different applications electric reflecting many views and usages of information.

The model must also facilitate implementation by hundreds or even thousands of diverse system implementers. This is a significantly different situation for most IT system developments. To achieve these goals, a model is required with the following characteristics:

- It must be a formal expression of the entities identified in medical information. A formal model can be validated, is implementable, and enables conformance testing (for example, do our databases and applications conform to GEHR version xx)
- It must facilitate evolution (rather than revolution) in existing system implementations and data as well as the construction of new implementations. It is important that any proposed model and exchange format adopted does not leave large numbers of existing systems 'out in the cold', unable to develop towards compliance with the proposed standard EHCR architecture.

To satisfy these requirements (and many others) a pure object-oriented modelling formalism has been chosen. Some of the advantages include:

- An OO formalism has a superset of the semantics found in other formalisms, such as encapsulation of data and behaviour, modelling of incomplete concepts, and inheritance; it is therefore capable of expressing existing concepts as well as more powerful ones;
- Since the primary construct is the class, which can address the model of real- world entities, it is directly comprehensible to human beings. It is also of course comprehensible to a computer

by means of expression in an OO programming language and the use of programming tools;

- Rigorously defined ‘views’ of the model can be created for less semantically powerful but extremely important non-OO technologies such as relational databases, and exchange mechanisms such as ASN.1

Inheritance is a major feature of an Object-Oriented formalism and can be used in very powerful ways to formally define and enforce relationships between different levels of the model, and between the model and OO implementation. For instance, an implementer can directly create a subclass of the ‘EHCR’ class found in the model and guarantee conformance (enforced by the compiler) while being free to expand the idea of EHCR to fit that existing at the site. Furthermore, it facilitates the creation, understanding and maintenance of non-OO views such as relational database table definitions.

It should be re-emphasized that use by the GEHR project of this formalism does not imply that GEHR compliant record systems would be required to implement all its features. Neither does it imply recommendation for use of object-oriented database methods. The GEHR object model is intended to capture the full set of GEHR care requirements which affect data structure and communication. Good ethico-legal practice must be supported by the medical record, but its accomplishment also rests on sound procedures beyond the remit of GEHR. Legislation in this area may influence clinical procedures and may imply constraints on the data structures and organization within the record.

The GEHR Object Model (GEHR OM)

Presented here is an abridged part of the GEHR OM expressed in the Rumbaugh notation [see Figure 8.24]; the diagram is derived formally (and in future, automatically) from the equivalent formal textual definition expressed wholly in the Eiffel language. GEHR is not wedded to a particular modelling notation, but to the use of a formalism for which a public domain definition and tools exist.

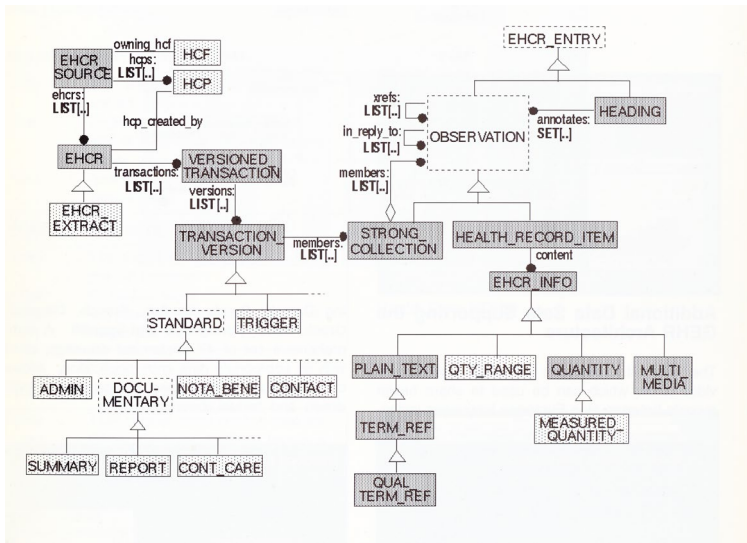


Fig. 8.24 The earliest formulation of the GEHR Object Model in 1994. In subsequent years, it was separated into the Synex reference and clinical data object model and aligned with the Ocean reference and clinical data archetype model, to provide the foundational openEHR architecture, in 2002. Image created by the GEHR project team (1994), CC BY-NC.

While there are many modelling ‘formalisms’ available, only a few have true formal definitions and are powerful enough to model abstract concepts. Many popular OO diagramming notations in use today have no formal definition, and therefore no reliable way of creating implementations and other views of the model without a lot of specialist human effort.

The formulas used by GEHR (currently the Eiffel language) avoids these problems, while providing powerful modelling and implementation capabilities, as well as tools.

The GEHR Exchange Format (GEF)

When exchanging EHCR data between sites, it is essential to preserve the structure and meaning of the data, while recognizing that different platforms, databases, and languages may be in use at each site.

A GEHR Exchange Format (GEF) has been developed which is independent of these site-specific aspects [see Figure 8.25]. It is designed to work if the two sides can represent their EHCR data in a way which conforms to the GEHR object model.




Fig. 8.25 Early ideas for a technology and vendor neutral care record exchange between systems. Image created by the GEHR project team (1994), CC BY-NC.

The Abstract Syntax Notation (ASN.1) has been chosen for the GEF since it is able to describe complex data objects which are all derived from a set of elementary types formally defined in the official ASN.1 recommendation CCITT X.208 / ISO 8824. The basic encoding rules of ASN.1 can directly be used as described in CCITT X.209 / ISO 8825 to produce a formal transfer syntax.

ASN.1 can express and encode the fact that Health Record Items (and HRI Collections) are *contained* in Collections, themselves contained in Transactions, according to some original structure.

The basic ASN.1 rules provide a simple mechanism using 'tags' for encoding all the various structural components of the EHCR. Tags have specific coding rules, so that any EHCR structure will be clearly identified in the exchange format.

Additional data sets supporting the GEHR architecture

The termset, comprising over 2000 entries, provides terms which can be used to share health records between nine European languages including Danish, Dutch, English, French, German, Greek, Italian, Portuguese, and Spanish. A comprehensive set of 47 anatomical drawings, used with an appropriate data entry application, allows clinicians to annotate outlined sites on the body with drawn and textual observations.

Specifications for GEHR data access and integration tools

The project has also built and provides specifications for tools to interface and integrate the GEHR architectural formalism (exchange format and object model with the emerging range of standard messages and architectures from specialized domains and systems. For example, laboratory data (OPEN-LABS), images (DICOM), ECG (SCPECG), prescribing (OPADE), systems (HELIOS), and such generic messages as HL7 and communication standards such as EDI.

Some GEHR-compliant health care records software products and prototypes

These prototypes were implemented at test sites in Belgium, France, Luxembourg, Portugal, London, and Hull/Sheffield.

The immediate future objectives of the GEHR Consortium

Raise awareness of GEHR as providing a coherent basis for the kernel of a common architecture for the contents of a comprehensive multimedia health record.

- Further develop the architecture, with an associated object dictionary towards an umbrella architecture for medical information across projects and clinical domains.
- Disseminate the Deliverables of GEHR in the public domain and participate in the future work of AIM and CEN towards a common architecture and standard for the EHCR.
- Disseminate GEHR object model, exchange format and tools specification in the public domain.
- Explore models of the EHCR created in specialist clinical domains and map these to GEHR object model, query, and exchange format.
- Explore interface of knowledge-based systems frameworks with GEHR object model, query and exchange format
- Explore emerging health care guidelines and map these to GEHR object model, query and exchange format.
- Establish mechanisms for generation, evaluation, and standardization, in appropriate domains of the fundamental GEHR record constructs: transactions, items, collections, context attributes.

- Work collaboratively with other clinical and technical groups and projects to establish a set of user test sites/demonstrators where common architecture compliance is specified, implemented, and evaluated.

Deliverables

4. GEHR requirements for clinical comprehensiveness
5. GEHR requirements for portability
6. GEHR requirements for communication
7. Specification of functional requirements for clinical use
8. Ethical and legal requirements of GEHR architecture and systems
9. Education requirements of GEHR architecture and systems
10. GEHR general syntax and semantics - interim report
11. GEHR implementation software tools - interim specification
12. Final systems report with evaluation of architecture and tools
13. Final clinical report with evaluation of architecture and tools
14. Final architecture report: details of specification and maintenance
18. GEHR software tools: final specification
19. GEHR final description
23. Documentation and maintenance procedures
24. GEHR users' manual

Preservation of meaning and clear delineation of the boundary of the record, in terms of authorship, access and accountability, were central to the GEHR architecture clinical functional specification. The technical architecture and its associated information model are further discussed below in the more detailed context of development of the work over the next ten years, leading to the creation and delivery of the methodology of openEHR.

Now, thirty years on, the original GEHR architecture has very substantially evolved. In continued fulfilment of its founding principles, it is openly specified, instantiated and widely implemented, in different software technologies and by different vendors of systems, under the aegis, now, of the openEHR Foundation and the openEHR International self-governing Community Interest Company (CIC). The need now is for a similarly cohesive and concise principled framing of an architecture for a care information utility, centred on this architecture of the care record.

openEHR

At the time we were creating the openEHR Foundation, I was working closely with UCL Medical School deans and heads of departments, to pull together the wide-ranging community of clinical researchers, creating a database of all investigators and their projects. I was also responsible for coordinating the merger of disparate IT support teams right across the clinical and life science faculties of UCL and its merging medical schools and research institutes. I was doing all this as a member of the UCL Biomedicine management executive, and thus closely in daily contact with its leaders and with the wider university management, also as a member of its Information Strategy Committee and later as a member of its Finance Committee. These were important integrative academic roles, strongly dependent on IT, and thus the sorts of things I was appointed for and expected to take on. They gave me a wider position of status and trust, and thereby some shelter for the highly vulnerable nature of a pioneering mission like health informatics within a medical school. I was asked to chair the UCL-wide Infrastructure Committee, overseeing the changing relationship of corporate academic services and academic departments across all faculties. In the context of the clinical mission, I represented the University in its research linkages with and handling of clinical data arising in the everyday health care services of its partner NHS trusts, and on the groups drawing together the IT teams of each trust, through local mergers and implementation of the NHS National Programme for IT. I was also asked by the NHS to create and lead a national academic forum for health informatics, bringing together people from all UK universities.

My UCL academic department thus became a hub and hive of connections across academic and professional, technical, clinical, organizational and public and private sector bodies. It connected with students in many faculties and in the wider NHS, on many levels. Leading all this was a complex, multi-faceted and busy enterprise! My staff were resilient folk and took on the challenge of finding their way through what was often a jungle in those anarchic years of transition in health care and information technology, much as I had done in my early years at Bart's. As it had been for me at Bart's, this was a tough ask, especially so being positioned in a medical school environment. Such environments can often feel a bit like the Wild West, but they come with the compensating reward of independence and freedom to explore, which are essential in all creative endeavours.

It was around this protected, both interdisciplinary and multiprofessional base, that the mission of openEHR was created and evolved. It did so with members and external colleagues of the department who already had, or went on to achieve, eminent positions of academic and professional

leadership elsewhere in the country and the world. All drew on and contributed immeasurably to the rich and complementary connections of different perspectives and endeavours made possible within this unique environment, as I further discuss in Chapter Nine.

From the foundations of the GEHR project, the progression to the openEHR architecture and its associated curated archive of clinical data models of today has been a lengthy, challenging and often perilously insecure story of iterative and incremental implementation and adoption. It has involved navigating changing scope and requirements, architectural refinement, implementation and testing, team and organizational development and, most crucially, growing industry, health care organization and governmental adoption and partnership.

From the original GEHR object model has evolved the openEHR reference model of today, which includes well-tested and stable classes of data descriptive of the who did what, when, where, how and why, of digital care records. And from it and subsequent projects and products have come iterated versions of the archetype constraint model that overarches the models of clinical data captured, searched and communicated, in and through these records. A new scope of decision support and work-planning has been pioneered, notably by Rong Chen at Cambio Healthcare Systems and Thomas Beale at Ocean Informatics, Ars Semantica and now Graphite Health.

Open-source specifications, tooling and other shared code have been contributed by partners in the growing international openEHR community. The openEHR website is the gold standard of recognition of these. Those with which I have more closely connected include: Ocean Health Systems (Archetype 1.4 Designer, Clinical Knowledge Manager), Cambio Healthcare Systems (early openEHR Reference Model classes, open-source) Better Healthcare (Archetype 2.0 Designer), Thomas Beale (openEHR Eiffel workbench, lead curator with Sebastian Iancu (Code24) of the openEHR specifications and their technical governance, basic metamodel and expression language), NEDAP (Archie 2.0 reference model classes, open-source), Seref Arikan, UCL/David Ingram (Opereffa openEHR platform, open-source), Pablo Pazos/Cabolabs (EHRServer openEHR platform, open-source), Tony Shannon and Christian Chevalier/Ripple Foundation (EtherCIS openEHR platform), Rob Tweed/M/Gateway Developments (QEWD openEHR platform tooling, open-source), Vita Group (EHRBase openEHR platform, open-source). Pablo Pazos and Vita Group are likewise working on a software framework to test and accredit software products for their practical operational conformance with the openEHR specifications.

Through the foundational work and collaboration of Heather Leslie, Silje Ljosland Bakke and Ian McNicoll, and many others, now being led by Paul

Miller, Joost Holslag and Vanessa Pereira, the body of clinical models has grown to contain the largest curated set of datapoints of any such resource in the world. This has been a phenomenal achievement and the translations of these archetypes into multiple languages, with their different alphabets is, in itself, a highly significant offering that is openly licensed to the whole world. Again, the openEHR website is the gold standard of recognition of these contributions. Hanna Pohjonen and Heather Grain are likewise bringing new focus to openEHR educational resources and accreditation of course providers. And a first cohort of openEHR ambassadors and openEHR affiliate organizations has evolved to represent openEHR mission and localize its adoption in different national jurisdictional contexts. This again has entailed learning by doing: we have usually imagined these entities and allowed time to pass before formalizing them legally, in their scope and relationship to the main openEHR International board.⁸

This pathway has been long and meandering, all the while seeking to steer a straight course along crooked lines. Rather than rewrite more detail of the history, here, and risk unintended errors of omission and commission, I have assembled key documents that I have written over time and include them as a set of Annexes to this half chapter, in the online archive of additional resources of the book.⁹ These trace the origins and development of the openEHR mission, alongside documents recording and acknowledging foundational contributions and contributors. They provide an important record, and I will aim to continue to update these to ensure they remain consistent, inclusive and supported statements. The online and freely accessible archives of the openEHR specifications and related clinical models acknowledge their multiple contributors very carefully, from a worldwide community that crosses disciplines and professions.

Annex I, written in 1999, ten years after the commencement of the GEHR project, expresses the mission and rationale of openEHR. Annex II describes the origins of openEHR up until the establishment of the openEHR Foundation, in 2002. Annex III is a transcript of my lecture at Medinfo 2007 in Brisbane, which is also on *YouTube* and accessible from the openEHR website. The history of the movement from 2002 until the creation of the Community Interest Company, openEHR International, in 2018, is

8 As I finalize this book's manuscript, after its copy-editing by Open Book Publishers, Sam Heard has just taken on the chairmanship of the openEHR Foundation from me, and Rachel Dunscombe has been appointed to be the first Chief Executive Officer of openEHR International. It is a wonderful time of transition in openEHR mission, hugely enabled by the openEHR International Board.

9 Available at <https://www.openbookpublishers.com/books/10.11647/obp.0335#resources>

recorded in Annex IV, and the updated 2020 vision and mission, written with Thomas Beale, is at Annex V. An updating account of the current status of the movement will be added to, from time to time, which I will ask new leaders of the mission to help create, and perhaps then take over.

openEHR is now a proven vendor and technology neutral architecture for digital health care records. Through its decades-long evolution and stabilization, it has acquired much new structure and form which is best not paraphrased here but taken up through study of the website. It now embodies clinician-defined models of record content (the models with ISO-standardized generic structure, known as openEHR archetypes) linked with internationally standardized clinical terminologies and translated into many languages. The reference model of generic building blocks and associated software tools enable non-technically literate health care users to design and specify a care record. By design, this record can then be hosted as an application on an openEHR standardized and generic technical platform for health computing. This is akin to an Android platform for health care with the health record running as an Android App. There can be many such applications, supporting the many specializations, workflows and connections of health care services, but they share a common semantic core, embodied in the reference model and archetype model.

The foundational ideas about record architecture embodied in GEHR through classes of Health Record Item, Collection and Transaction, have iterated, metamorphosed and extended to element, composition, cluster, entry and folder, all strictly version-controlled, and connected into a model of recorded observation, evaluation, instruction and action, and their combination in workflows and timeseries of events. The archetype model and expression language have widened in scope to enable close alignment and binding with evolving clinical coding, classification and messaging systems, such as SNOMED (Systematized Nomenclature of Medicine), ICD (International Classification of Diseases), Logical Observation, Identifiers, Names and Codes (LOINC) and Fast Healthcare Interoperability Resources (FHIR). These have, perhaps, reached tractable limits of relevant elaboration for general purposes, and Occam's razor looms. They will likely continue to evolve through simplification of essence and refined granularity of description. The Ocean Health Systems Clinical Knowledge Manager (CKM) tooling has advanced and now hosts substantial libraries of both nationally and internationally curated clinical data models. The Better Care and Ocean company tooling for openEHR archetype design are used widely around the world, as is the Vita Group EHRBase open-source openEHR platform implementation.

Conformance to this open platform specification ensures that data can (with due and informed care!) be reliably queried across all component

software applications in a standardized manner. It holds the promise of loosening the harmful hold of proprietary data formats, that currently prevail in health IT systems. By adopting the openEHR platform approach, users can build a local set of inter-compatible systems, choosing from the products and services of a growing community of both large and small providers, or developing their own. Most importantly, the transparent clinical and technical discipline of openEHR provides robust ethico-legal foundations on which to satisfy growing international concerns and expectations for confidential handling of personal data, as exemplified in the general data protection regulation (GDPR) in Europe, for example. Huge sums are spent in health informatics, but very much greater value will be achievable when shared methods and resources prevail, as they do elsewhere in science and engineering.

A key focus of new learning is about implementation reaching towards the vision of a citizen-centred care information utility. Leadership in this quest will be centred on pioneering openEHR adoption in various care sites, involving whole health care organizations and regions and their supporting industries, in new partnerships. A most encouraging development has been the adoption of openEHR as a standard for care record repositories in substantial clinical initiatives of health care providers, such as cancer services around The Christie Hospital in the Wirral in England, the care record ecosystem of the Karolinska University Hospital in Stockholm, and the health care system of the state of Catalonia, in Spain. Better's OneLondon project has been another and widely acclaimed advance along these lines. There is also growing confidence demonstrated by pioneers in India, and in the activities of openEHR ambassadors, national affiliate organizations, and the now more formally established clinical and educational programmes, which are all starting to gel in a complementary fashion, under the increasingly securely-established CIC board, ably co-chaired by Tomaž Gornik and Rachel Dunscombe. Jordi Piera-Jiménez, leading for the trail-blazing adoption of openEHR in Catalonia, has provided notable leadership and advocacy in world fora. Details of all this can be found on the openEHR website.

Health care information systems have become opaque and entangled – new discipline is needed to unravel them

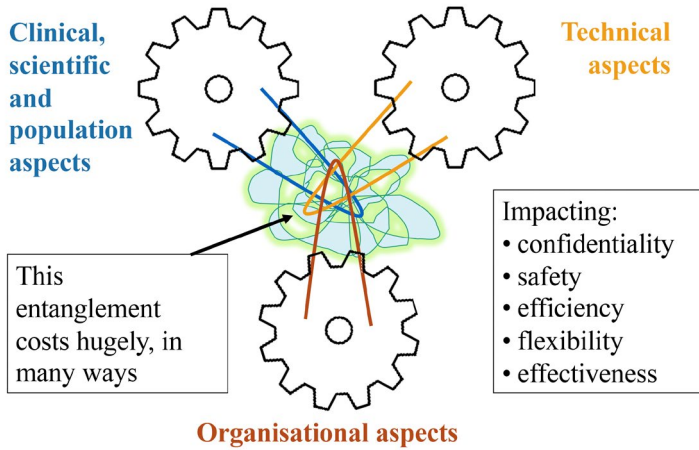


Fig. 8.26 The entanglement of non-coherent information systems. Image created by David Ingram (2010), CC BY-NC.

openEHR was created to help disentangle clinical systems, so they can interoperate rigorously and meaningfully




Fig. 8.27 Interoperability achieved through the common ground of a coherent semantic framework. Image created by David Ingram (2010), CC BY-NC.

I often used these two slides (see Figure 8.26 and Figure 8.27) to illustrate the entanglements of clinical data that can arise when they are interchanged between systems, devoid of a unifying semantic framework of meanings, and then illustrate how we are working to untangle the knots, using the common ground of openEHR. The cogs are pictured properly, here, reflecting the engineering rigour we have prioritized in openEHR, and they are revolving in synchrony, in contrast with Figure 5.1 and its imagineered depiction, which is actually of a non-functional, broken machine! I focused most on providing and sustaining an overarching vision and inclusive team environment, to hold and keep mission and people together. I was fortunate beyond imagination to have been a leader among a team of great leaders and have their commitment, support and trust.

From Opereffa to EHRBase—Towards a Standardized openEHR Open-source Platform

As wider interest in openEHR started to grow from the early 2000s, our previously closely knit team, centred on UCL, spread and expanded along different paths. Ocean and Cambio were early pioneers, and Rong Chen at Cambio created a set of Java classes to implement the openEHR reference model and make these available open-source, to help others experiment with the creation of openEHR-based systems. In parallel, my colleagues, Dipak Kalra and David Lloyd, working at UCL, switched their focus to leadership of a CEN project team established under a Working Group chaired by Gerard Frericks, to work on a standard architecture for electronic health record communication. This involved them in several years of intensive work, subsequently adopted in CEN and ISO as the 13606 standards. In its two releases, this incorporated the openEHR archetype model as a public domain standard for the clinical content models used to specify content of the care record shared between systems.

In 2008, Seref Arikan came to London as my doctoral student, to build on his extensive experience, both as an economist and computer software developer for hospital systems in Turkey. This connection evolved into a most fruitful and consequential professional partnership, supported jointly between CHIME and Ocean, and from it have come new insights that have strongly influenced implementation. In his PhD dissertation, Seref explored how Bayesian networks could be integrated with the architecture of openEHR. He invented a new software paradigm and model for implementation of queries into structured care record data, unravelling and resolving ambiguities that had arisen in the overlap of terminology used to

describe clinical knowledge and the information model used to structure and organize it within a clinical record database.

In those early years, I was concerned about the legal integrity of openEHR intellectual property rights. Through my involvement in the StartHere community and with the invaluable help of its leader, Sarah Hamilton-Fairley, and the chief executive of UCL Business, Cengiz Tarhan, I was introduced to and commissioned Oliver Bray, a partner specializing in intellectual property rights at the leading London law firm, Reynolds Porter Chamberlain, to draft documents formally assigning the openEHR intellectual property rights (IPR) to the by then established openEHR Foundation, from its founding member organizations, UCL and Ocean Informatics.

I was further concerned that we make a reference implementation of openEHR available, as an open-source platform (akin to an Android or Unix of healthcare), to facilitate dissemination and uptake. This was tricky territory as early pioneering companies partnering us were, very understandably, anxious to maintain their advantage as early adopters of the openEHR methodology, in their products and businesses. These conflicting interests required careful and sensitive balancing, in the interests of sustaining our partnerships and finding common ground on which to pursue the purposes and goals to which we and the openEHR Foundation were jointly committed and dedicated.

With Seref, I funded and pursued an early project to create such an implementation. He christened it Opereffa. Tony Shannon was likewise a committed advocate for open-source health care software and contributed his time generously to our project. Here is the diagram (see Figure 8.28) that set its scope and context. It was published on GitHub and the Opereffa software was downloaded and referenced in projects in many countries. It invited but attracted no more widely contributed code. This was a salutary lesson that open-source can be seen as a useful bootstrap by others—a free good—but not one they are always committed to or able to advance. The Opereffa initiative and its pioneering first steps towards an open-source platform implementation of openEHR, now being realized in the EHRBase project in Germany and the Cabolabs EHRServer project in Uruguay are described in a 2013 paper by Ingram and Arikan.¹⁰

10 D. Ingram and S. S. Arikan, 'The Evolving Role of Open Source Software in Medicine and Health Services', *Technology Innovation Management Review*, 3.1 (2013), 32–39, https://timreview.ca/sites/default/files/article_PDF/IngramArikan_TIMReview_January2013.pdf

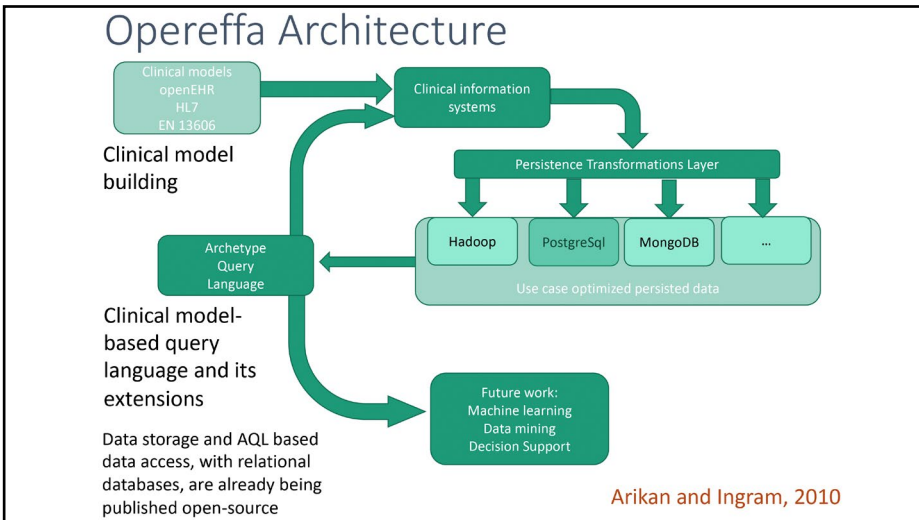


Fig. 8.28 The scope of the first open-source openEHR platform, Opereffa. Image created by Seref Arikan and David Ingram (2010), CC BY-NC.



Following from his engagement with Opereffa, Tony Shannon joined forces with Christian Chevalier, from Cambodia, who had previously visited me at UCL, along with a French clinical colleague interested in making a professional and commercial link with openEHR. Together they pioneered the EtherCIS platform, later joining up with Jo Milan's former colleague, Rob Tweed, who added his invaluable Marsden background in the Massachusetts General Hospital Utility Multi-Programming System (MUMPS) language. Rob capitalized on this to write a developer toolkit (QEWD) for openEHR compatible implementations under the Ubuntu operating system, including on the cheapest of microcomputers, the Raspberry Pi. The team had local successes in implementing a prototype shared care record in Leeds, but the battle to secure longer-term funding defeated them, despite Tony's Herculean and very public-spirited efforts over more than a decade. He eventually moved back to Dublin and took on a key clinical role overseeing policy for health care IT in Ireland.

In about 2015, Vita Group, a German company connected, I gathered, with the group that developed the worldwide success story of a prominent business software suite, entered the arena. Tony organized a meeting with them in London, in 2018 and, over time, the EtherCIS project metamorphosed into a solidly anchored new development activity, to

bring an Apache open-source-licensed openEHR platform to the level of a rigorous and dependable product. This has drawn on the already wide-ranging expertise of Vita Group and brought Christian Chevalier and Pablo Pazos into their development team, alongside Birger Haarbrandt and Stefan Schraps. Pump-priming resource was received from the German government in the HiGHmed project, tasked to connect clinical data and records from principal German academic medical centres onto an openEHR platform, linked with additional tooling to enable integration with HL7 FHIR messages communicated between systems.

This open-source initiative was christened EHRBase. It is the long-sought realization of an openEHR reference implementation, some fifteen years after my early championing of the importance of such an implementation, in growing and disseminating openEHR methodology and community. It is a stepping-stone that enables new entrants to the field to implement innovative clinical projects and applications, quickly, agilely, credibly and scalably. (Microsoft Word doesn't like 'scalably', suggesting I mean 'saleably'—I'm sure that is true, too!) Much to my pleasure, this is starting to happen widely across the developing world.

For many decades, such innovation has been impeded and prevented by the start-up cost and effort involved in getting a new, closed source and proprietary platform up and running. The USA Veterans Health Information Systems Technology and Architecture (VISTA) initiative, with its considerable backing, struggled to sustain and update such an infrastructure, focused on health-system-wide application. Such endeavour typically requires many years of effort, building and sustaining a substantial software team and guided by only a small fraction of the international clinical community input that openEHR has mobilized. openEHR clinical models now constitute the world's largest corpus of tried and tested, formally assured, openly licensed care record content models. These models are developed and live within their local domain of clinical requirement, service and governance, while able, as required, to interoperate within the widening range of health care systems that have adopted the standardized openEHR care record platform architecture. The Ocean Health Systems Clinical Knowledge Manager software is coming into its own as a tool for curating these knowledge models in context of local community needs.

The incremental emergence and backing for EHRBase, is transforming the market alongside the now powerful openEHR-based products of early closed source offerings of companies like DIPS in Norway, Better Care in Slovenia, Cambio in Sweden and Ocean Health Systems in Australia, with other well-established companies now starting to join in.

openEHR and OpenEyes

OpenEyes has a shorter pedigree than openEHR and arose as a mission and preoccupation for me after meeting Bill Aylward, its founder. I introduced him as one of the pioneers whose contributions were described and celebrated earlier in Chapter Eight. I have described there and elsewhere in the book how these two missions connected along my songline.

I met Bill and connected with OpenEyes in the early stages of its prototype development at Moorfields Eye Hospital. Bill had clout and credibility there and it underwrote his team. He had free rein and ability, personally, to link and oversee both the clinical purposes to be served and technical prototyping of the software. We worked together closely from its early days. He joined the openEHR Foundation Board and I the board of OpenEyes. Had openEHR then existed in the form of what is now the EHRBase platform, OpenEyes would have been built on that platform from the start. Bill, Seref Arikan, Ian McNicoll, Mark Leaning and I secured funding and worked on building bridges between the two—for example, in designing an archetype for recording visual acuity and creating a software interface with high street ophthalmology services, working with the Black Pear software company.

To inform his PhD project, Seref and I worked alongside Bill to study and extract data from the legacy patient record database at Moorfields and take first steps to harmonize it within the framework of openEHR methodology. The tangle revealed in the underlying structure of the database of clinical records of this justly world-renowned institution was a shock. The local Trust IT team must necessarily have been substantially engaged in sustaining arcane methods for reliably feeding and interrogating it. Bill's motivation for creating OpenEyes, to improve on this underlying costly and unsustainable disorder, was well justified.

One great personal regret, that it is important to record, here, is that in mid-2013, Bill and I tried very hard to interest the National Health Service (NHS) in helping us integrate OpenEyes onto the evolving platform of openEHR. Sadly, this came to nothing. Our detailed proposal to bring together key initiatives in open-source clinical applications software of the time onto an openEHR integrated care record platform, that we worked on with Ian, Seref, Mark, Tony, Thomas and several other groups fell afoul of funding gremlins, somewhere. I saw several hundred proposals adjudicated under this NHS Integrated Care Records initiative, and, under the bonnet, few were much more than a Web-based patching together of separate applications.

Our ORSINI (Open Record Standards INitiative) vision, as we christened it, remains on the 'to-do list' for OpenEyes, as soon as possible, when we can

find a funder. This will, hopefully, soon succeed, as we gain implementation experience in the OpenEyes national programmes for Wales and Scotland, and alongside the groundswell of international usage and interest in the open platform approach. The Apperta Foundation has pioneered a new business model for open platform-based health applications. This has thus far proved viable, against the doubts and concerted opposition of some. It remains to be seen whether it can prove viable for attracting the level of new investment needed to align the current OpenEyes software onto an openEHR standardized platform, such as EHRBase.

In my new retirement roles at a distance from the more onerous fray, I see a pathway opening for our long-ago ORSINI platform to become a global utility, safely and sustainably, in a way that achieves much more, much faster and more cheaply. That has been my provocative ten-ten-ten mantra (ten times in each of these aspects) of recent years, along with the buzzphrase of 'Small' or 'Little Data', which I alighted on to focus attention on the ground-level data captured using the 'omniscles' (see Chapter Three) of openEHR archetypes, as a necessary foundation of less noisy Big Data.

openEHR and OpenEyes are both landmark initiatives that started from ones, twos and threes of participants, grew to tens and hundreds, and are now, in a similar leap, moving towards the tens of thousands. Such transitions are similar in terms of the energy, staying power and persistence they require. This kind of experience has informed my approach in leading disruptive endeavour, holding together teamwork, culture and environment at the centre of connected mission. As activities in my legacy portfolio, they are chicken and egg-like, now coming increasingly into focus as a unifying thread of initiatives at national levels. In terms of their specifications and implementations, they demonstrate a modular architecture of clinical record applications, and a modular platform architecture to house and interconnect families of such applications. OpenEyes is now (albeit very slowly because of resource constraint) on a pathway of migration to fully align its data with the archetype model standard of openEHR, following on from our initial foray in capturing visual acuity data in this way, some years back. It is disappointing that health informatics policy in the home country of openEHR and OpenEyes has somewhat lacked acuity of vision and been somewhat short-sighted in its goals. Perhaps this an inevitable reflection of its huge size. Scotland and Wales have established teams to work on a national openEHR platform and ecosystem of open-source clinical applications.

In parallel, OPENeP for medicines management (a significant missed opportunity for the NHS in not heeding Apperta's attempt to secure this, too, in the public domain, I gathered) has been developed and marketed by Better Care and is largely already constructed and operating within its

openEHR-based platform framework. New applications to support routine clinical ward level observations of vital signs, linking with new integrated instrumentation for their capture (open-eObs), and patient reported outcome measures (openOutcomes), using internationally standardized questionnaires, are advancing rapidly within the Apperta Foundation, and creating considerable worldwide interest. These are all building blocks of future care information utility.

Why has all this failed for so long to crystallize in policy circles? My former student and colleague, Dipak Kalra, who now works as a consultant at senior industry and government levels across Europe, told me in a March 2023 e-mail that the message is at last beginning to take root at those levels. As purpose and goal, the concepts and their early implementations were articulated and visible, twenty years ago, and presented at the highest intergovernmental levels, as many stories in the book describe and bear out. The legacy of frequent new initiatives and poorly interoperable systems and applications *in situ*, that health care IT teams have had to adapt to, choose from, and keep running, have kept them distracted by short-term issues, with few able and equipped to take a longer-term view.

More detailed history, spanning now thirty years of openEHR and twenty years of OpenEyes, can readily be traced online. I retain the most detailed personal archive of those decades, charting there the progress along my songline. I do not wish to paraphrase or rewrite them. Stuff happens and things change but it is not for me to revise the history of something I was so closely connected with. Others must now traverse their own songlines, tell their own stories, and write their own histories.

Through this reflective half chapter, I hope the reader will find anthro-vision—Gillian Tett’s new term¹¹—and insight into the human side of openEHR and OpenEyes, as this is where their sustained success had its origins and retains its staying power. It is a hard story to tell as it has been very hard to achieve, and hard to survive, for many people. Of course, there has at times been contentious, manipulative and angry dispute and disappointment. I have not written about these instances, to preserve confidentiality about hard times that many have been through, and still do, and will likely prefer not to be reminded of! Learning from experience of this human dimension is where important lessons for the future of care information utility will lie. Environment is all important—a place and setting where people are enabled to connect, work and be valued, both collectively and independently. Teams, and their complementary perspectives, goals, roles and motivations, are all important. Connecting with and capturing opportunity, and using it flexibly

11 G. Tett, *Anthro-Vision: A New Way to See in Business and Life* (New York: Simon and Schuster, 2021).

and well, are paramount. Above all, what have mattered most have been trust and staying power. What matters in anarchic transitions is to cope and get on with the work, see it through and preserve important values and meanings. Otherwise, much more can and will be lost.

Bill and I have both stepped aside from our central roles in OpenEyes and openEHR. Me to dance with my wife, look after our health, enjoy family, freedom, fun and opportunity to branch out in new directions, and write this book! And Bill, a decade or so younger than me, to sail and enjoy the world with his wife. Stepping down is as important as stepping up. Innovation does not happen unless people step up to carry the load. It is not sustained unless they step down, too. That is the way in which current postscript is translated into future preface, and a new generation inherits and uses the useful legacy landmarks and stepping-stones left to them, and learns from the less advantageous ones, as well, as health care continues to evolve and move forward. Bill stepped down and James Morgan, Peter Coates and David Haider have powered his vision of OpenEyes forward. As the openEHR founders have gradually stepped down, a new board and team of complementary leaders, brimming with energy and commitment and reaching to wider horizons, is powering their vision forward, too. As a seasoned grandfather, like myself, a principal reward in life is to have confidence and pride in your family. openEHR and OpenEyes are also great families and I do in them, too—with head still, just about, and heart, always warm, but not so much with hand, other than to clap!

Parenthesis—And So?

This half chapter has described long-term work dedicated to imagining, developing and implementing tested solutions for a central component of the care information utility—a globally standardized, openly specified and freely accessible methodology supporting a locally governed, citizen-centred platform ecosystem of digital records of care for whole health economies. In parallel, and in connection with the profile of its founder, Bill Aylward, in Chapter Eight, it has connected with his ambitious and highly successful quest to create a clinical community motivated and led, widely adopted and sustainable, open-source medical record for eye care. These interdisciplinary, multiprofessional and multi-sectoral initiatives have explored practical issues of implementation, in detail and at scale, and the creation of new community and legal organization to lead and coordinate their further evolution, as viable and cost-effective community interest endeavours. And so, where next? That is for others to envision and lead. It will be what they now create from the legacy that they inherit. How they do it will matter as much as what they do.

