

It's Time to Replace Steam Engine EMRs with AI-EMRs

David Scheinker, Stanford University School of Engineering, Clinical Excellence Research Center, Stanford University School of Medicine; and Matt Hollingsworth, Recombinate Health

Contact: dscheink@stanford.edu

Abstract

What is the message? Historically, replacing one large steam engine with one large electrical engine yielded minimal productivity gains. The industrial revolution came with the redesign of the production process. Modern Electronic Medical Record (EMR) systems are like steam engines in the industrial era – outdated technology poorly suited to leverage the potential of AI. Modern hospitals need a new AI-EMR built from the ground up to fundamentally transform healthcare delivery, quality, and cost structures. The authors propose that major tech companies like Amazon and Google should partner to purchase a small hospital system to develop a new AI-native EMR.

What is the evidence? The authors cite multiple sources demonstrating how current EMR deployments limit the potential impact of AI. They provide examples of complementary technology, such as mathematical optimization or learning from randomization and experimentation, broadly used in other industries to harness the full potential of AI but incompatible with current EMRs. The paper contrasts healthcare's expensive, fragmented approach to technology with more efficient ecosystems like Apple and Android, and notes emerging market movements in this direction, including Oracle's \$28 billion acquisition of Cerner and venture capital investment in healthcare delivery systems targeting technological transformation.

Timeline: Submitted: April 10, 2025; accepted after review April 11, 2025.

Cite as: David Scheinker, Matt Hollingsworth. 2025. It's Time to Replace Steam Engine EMRs with AI-EMRs. Health Management, Policy and Innovation (www.HMPI.org). Volume 10, Issue 1.

The invention of the electric motor had relatively little immediate impact on productivity. Initially, companies swapped one large steam engine for one large electric engine. Productivity increased only after factories redesigned the production process in ways made possible by each machine being powered by its own electric motor.^(1,2) As Google continues to announce large language models (LLMs) that extend their lead over physicians in a growing set of clinical tasks,^(3,4) hospitals are investing in the modern day equivalent of installing light bulbs in a factory powered by one large steam engine, the Electronic Medical Record (EMR). To unlock the potential of AI to improve quality and reduce costs will require the development of a new AI-EMR. Here, we lay out the potential benefits by drawing on examples from Amazon, Apple, and Google and a roadmap for a partnership with the appropriate capital and expertise to do so.



we'll fix the formatting / text wrapping

Figure: A hospital figuratively powered by a steam engine EMR - one large, outdated system on which all the machinery within relies.

Potential Benefits

Large language models (LLMs) have drastically improved their ability to interpret simple medical text, aid physician decision-making, and tackle a variety of administrative tasks. Hospitals are racing to adopt LLMs as: digital scribes, a replacement for Google for medical questions,

administrative assistants, and to draft everything from responses to patient messages to prior authorization requests (3-6). The initial results are promising, but the cost, quality, and efficiency of these use cases is fundamentally handicapped by their dependence on the hospital's EMR steam engine.

An AI-first EMR would open the door to complementary optimization technology ubiquitous in other industries; facilitate the creation of a learning health system; improve quality and reliability; automate the low-value patient and provider tasks that consume the majority of time people spend interacting with the healthcare system; and lower rather than raise costs. Amazon demonstrates the potential efficiency allowed by AI and complementary technology. The Apple and Android ecosystems illustrate the potential of secure, connected data sharing. All of these companies improve productivity rapidly by building randomization into the consumer experience to learn from their own data.

In order to generate value, general purpose technologies such as LLMs require complementary infrastructure.(1,7,8) Mathematical optimization or mathematical programming is a central pillar of this complementary infrastructure in practically all large efficient industries. From early uses to solve logistics problems for the U.S. Air Force during World War II and to design pharmaceuticals in the 1960s, mathematical optimization is now widely used by retailers like Amazon and Walmart; all major airlines; electrical grids; financial investment firms; and even The National Football League.(9-11) To generate value from AI-based forecasts of consumer demand, retailers like Amazon and Walmart use mathematical optimization to create efficient shipping schedules for millions of packages per day that minimize delivery times and fuel costs.(12) To generate value from AI-based forecasts of weather patterns and potential delays, major airlines use mathematical optimization to assign planes to routes and crews to planes while minimizing fuel costs and respecting myriad practical and regulatory constraints.(13) Such optimization is part of the reason that airline productivity and costs have improved over the last 30 years while hospital productivity and costs have not.(14,15)

To create a schedule that respects myriad rules such as each team playing every team in their division twice and not flying across the country for consecutive games, the NFL uses mathematical optimization to create a draft that humans revise.(16) There have been hundreds of papers demonstrating how mathematical optimization can complement AI to improve surgical scheduling, nurse staffing, patient clinic appointments, and service-to-unit assignments, but,

because modern EMRs do not support mathematical optimization, there have been very few implementations.(17–20) A new AI-EMR would include this century-old complementary technology.

An AI-EMR could deploy functionality well-established by technology firms to inject randomization into the standard of care to generate evidence that helps the U.S. healthcare system to attain the long-sought goal of becoming a learning system.(21) For decades, technology firms have generated improvements from evidence collected by injecting randomization and experimentation into their technology and user experience.(22,23) They have leapfrogged healthcare, where experimentation and standards of evidence go back almost 3,000 years to King Nebuchadnezzar of Babylon. (24) Most clinical trials remain expensive, time consuming, and restricted to small populations with limited potential to generalize.(25) Healthcare has taken promising steps in this area, such as the inclusion of pseudo-randomization into the standard of care for algorithm-enabled remote patient monitoring of youth with diabetes, but as individual engineering efforts outside the EMR.(26) Progress would improve rapidly with an EMR engineered to allow randomization and experimentation as part of the standard of care.

Integrating all relevant information about a patient, to replace current context-free messages, would drastically improve quality and reliability. Recently developed model context protocol (MCP) technology connects LLMs to other software and databases in a standardized way that improves efficiency and reduces hallucinations. (27,28) Only with access to data can LLMs help nurses, physicians, or patients answer common questions like, “Which medications and doses were prescribed?” or “How much out-of-pocket spending will this entail?” or “How long was the ICU stay?” Only with access to data can a digital scribe ensure that it heard a physician in a noisy room say “hydroxyzine” not “hydralazine” or “.2 mg” not “2 mg.” On your phone, applications that auto-fill your passwords, measure your pulse, or let you talk to an LLM are possible only with secure, reliable connections to the FaceID, camera, and microphone.

Designing administrative and clinical workflows based on the current capabilities of LLMs, and projections about their improvement, would allow the elimination, rather than reduction, of numerous low-value tasks that drive burnout and reduce the time physicians and nurses have with their patients. Entire categories of non-clinical work currently done by highly trained people could be eliminated. Nurses could stop scheduling routine appointments and focus on patients

with complex needs. Physicians could save hours each day and eliminate pajama time by not documenting billing codes. Experienced nurses could review and approve AI-generated schedules or nurse-patient assignments rather than dedicate half their time to creating these themselves. Analysts could dedicate time to generating insights about opportunities to improve care or cut costs rather than putting together static dashboards. Apple and Android helped third-party travel applications improve convenience, choice, and competition by letting people compare the costs of rides to the airport, hotel stays, and plane tickets. The operating system makes this possible by providing the applications with location information, saved payment information, and a Bluetooth connection for the wireless hotel room key.

Perhaps the most pragmatic reasons are lower costs and better security. Current AI productivity improvements raise costs because current EMR deployments are customized to each hospital at the wrong level, e.g., to each hospital's preferences and needs at the "system level." Every application is tested by every hospital's security and stability teams, then it is customized and installed in an expensive 6- to 18-month-long process, and is often incompatible with user preferences. Each Apple and Google phone is customized to the user's preferences and needs at the "user level." Every app is tested by Google and Apple security and stability teams once, each user installs it with a click and customizes it to their preferences. The resulting process is more efficient and more secure. The fragmented nature of healthcare systems makes them a frequent target of large, costly attacks. It took dedicated teams of experts employed by the FBI months to break into an iPhone they had in their possession. CAN WE SWITCH THIS LAST SENTENCE WITH THE PREVIOUS ONE FOR CLARITY AND CONTEXT?

Table: Potential benefits and uses of AI-EMR technology with examples from other industries

| Potential benefits of AI-EMR | Additional technology built into EMR | Examples from other industries | Potential hospital uses |
|---|---|---|---|
| Improve procedure, physician, nurse, and staff scheduling | Mathematical optimization to complement generative AI for scheduling and logistics | Amazon, Walmart, airlines, and the NFL use optimization to schedule staff, supplies, and logistics | Schedule surgical procedures, nurse-patient assignments, and other aspects of operations using patient characteristics combined with physician and nurse experience and preferences |
| Generate evidence for a learning system | Pseudo-randomization and experimentation as part of the standard of care | Amazon, Apple, Google, and other technology firms generate data by integrating randomization and experimentation into the user experience | Enable virtual randomized clinical trials as part of the standard of care by, for example, allowing randomization in digital health interventions such as remote patient monitoring |
| Improve LLM clinical decision-support | Connections, such as Model Context Protocols, between LLMs and patient history and data | LLMs properly integrated into the programming infrastructure significantly improve the speed and efficiency of software development | LLMs provide more accurate, relevant answers to care providers by drawing on patient history and data |
| Fully automate tasks carried out by LLMs | Connections, such as Model Context Protocols, between LLMs and EMR functionality | Broad adoption of full-service chatbots for low-risk, low-value tasks such as scheduling or providing financial information | LLMs schedule non-urgent appointments and answer simple patient questions |
| Lower adoption and maintenance costs | A secure system of permissions for data sharing and functionality between the EMR and new applications. | Apple and Google application environments reduce the initial and ongoing cost to validate new applications and integrate them into the operating system | Low adoption and integration costs for direct access to a large ecosystem of applications such as clinical risk prediction, scheduling, or financial tools |

A Practical Roadmap

Modern patient care is too complex for a group of motivated innovators to build something in their garage. A company or partnership with expertise and capital, such as Amazon and Google, would have to buy a small hospital system; embed clinical, technical, financial, and operational experts in the care teams; and develop the new hospital AI-EMR from the ground up. There would be significant challenges. Fortunately, the resources necessary to overcome them are available, and the rewards of doing so are vast.

The market is already moving in this direction. Amazon has made several forays into the space, first through an ultimately unsuccessful partnership with other large companies to create a new firm, the creation of its own online pharmacy, and most recently with the \$4 billion dollar acquisition of primary care provider One Medical.⁽¹⁹⁻²¹⁾ However, these efforts are not yet targeted at hospitals nor the creation of a new EMR. Oracle, a large technology company, purchased Cerner, the second largest EMR after Epic, in 2021 for \$28 billion and is working to transform it into an AI-powered EMR.^(22,23) Unfortunately, that effort has run into the problems common to modernizing legacy technology. Venture capitalists (VC) have committed to the viability of a similar opportunity. A VC group recently finalized a \$485 million deal to acquire an 8,000-person integrated healthcare delivery system with the goal of using technology to improve its efficiency.^(24,25)

Purchasing a small hospital system would be a relatively small, low-risk investment for a partnership like Amazon and Google. Google recently acquired 1,800- person cyber security firm Wiz for \$32 billion, approximately 100 times their \$350 million annual revenue. In contrast, hospitals are typically acquired for several hundred million to one billion dollars, approximately one to three times their annual revenue and a small fraction of what Google spends on electricity each year. Purchasing a hospital system and building an EMR from the ground up may also be safer than purchasing an EMR and inheriting the problems Oracle is attempting to overcome.

Although the costs of transitioning to a new EMR are high, they could be defrayed by the massive short-term savings made possible by the AI technology Google has just developed and Amazon's expertise with efficient operations.^(3,4,12) An AI-EMR would embed or reduce the cost of the functionality on which hospitals currently spend hundreds of billions of dollars each year.

Because modern EMR's do not offer any kind of meaningful support for operations, practically every large hospital pays for systems to visualize their data such as Tableau, to schedule their employee such as Kronos, to manage their budgets such as EPSI, to deploy AI point solutions such as scribes, etc. In addition to the licensing fees, hospitals employ massive workforces dedicated entirely to supporting and maintaining the systems. The costs of the systems are high and the user experience poor, because companies compete based on their ability to sell to and integrate with hospitals, not on the quality of their technology. In contrast, markets with low barriers to entry such as Apple and Google app ecosystems see fierce competition to improve quality, user experience, and reduce cost. Finally, the value provided by these systems is limited by the difficulty, sometimes impossibility, of data sharing data. Nurse scheduling, the top operational expense of most hospitals, would be more efficient if it were powered by mathematical optimization integrated with data on: patient census, scheduled patient admissions, and detailed clinical data such as the specific needs of patients for whom only certain licensed nurses are appropriate.(26-28)

Modern patient care is too complex for a group of motivated innovators to build something in their garage. A company or partnership with expertise and capital, such as Amazon and Google, would have to buy a small hospital system; embed clinical, technical, financial, and operational experts in the care teams; and develop the new hospital AI-EMR from the ground up. There would be significant challenges. Fortunately, the resources necessary to overcome them are available, and the rewards of doing so are vast.

The market is already moving in this direction. Amazon has made several forays into the space, first through an ultimately unsuccessful partnership with other large companies to create a new firm, the creation of its own online pharmacy, and most recently with the \$4 billion dollar acquisition of primary care provider One Medical.(29-31) However, these efforts are not yet targeted at hospitals nor the creation of a new EMR. Oracle, a large technology company, purchased Cerner, the second largest EMR after Epic, in 2021 for \$28 billion and is working to transform it into an AI-powered EMR.(32,33) Unfortunately, that effort has run into the problems common to modernizing legacy technology. Venture Capitalists (VC) have committed to the viability of a similar opportunity. A VC group recently finalized a \$485 million deal to acquire an 8,000 person integrated healthcare delivery system with the goal of using technology to improve its efficiency.(34,35) As the

Purchasing a small hospital system would be a relatively small, low-risk investment for a partnership like Amazon & Google. Google just acquired an 1,800 person cyber security firm Wiz for \$32 billion, approximately 100 times their \$350 million annual revenue. In contrast, hospitals are typically acquired for several hundred million to one billion dollars, approximately 1-3 times their annual revenue and a small fraction of what Google spends on electricity each year. Purchasing a hospital system and building an EMR from the ground up may also be safer than purchasing an EMR and inheriting the problems Oracle is attempting to overcome.

Although the costs of transitioning to a new EMR are high, they could be defrayed by the massive short-term savings made possible by the AI technology that Google has developed and by Amazon's expertise with efficient operations.[\(3,4,12\)](#) An AI-EMR would embed or reduce the cost of the functionality on which hospitals currently spend hundreds of billions of dollars each year. Because modern EMRs do not offer any kind of meaningful support for operations, practically every large hospital pays for systems to visualize their data (e.g., Tableau), to schedule their employee (e.g., Kronos), to manage their budgets (e.g., EPSI), to deploy AI point solutions (e.g., scribes), etc. In addition to the licensing fees, hospitals employ massive workforces dedicated entirely to supporting and maintaining the systems. The costs of the systems are high and the user experience poor, because companies compete based on their ability to sell to and integrate with hospitals, not on the quality of their technology. In contrast, markets with low barriers to entry such as Apple and Google app ecosystems see fierce competition to improve quality, user experience, and reduce cost. Finally, the value provided by these systems is limited by the difficulty, sometimes impossibility, of data sharing data. Nurse scheduling, the top operational expense of most hospitals, would be more efficient if it were powered by mathematical optimization integrated with data on: patient census, scheduled patient admissions, and detailed clinical data such as the specific needs of patients for whom only certain licensed nurses are appropriate.[\(36-38\)](#)

Challenges

The primary risks of the proposed approach would be the clinical complexity of deploying a new EMR during clinical practice, the short-term declines in quality associated with the adoption of a new EMR, and the risk data blocking and legal action by incumbent EMRs.[\(39-41\)](#) Fortunately, the adoption of a certified EMR by practically every hospital in the United States has generated a vast literature of the common challenges and strategies to overcome them, as well as numerous

specialists with experience in multiple EMR deployments. The appropriate partnership could have, or could acquire, the expertise to meet these challenges. Amazon has a multi-decade record of successfully acquiring and improving the efficiency of low-tech, labor-intensive companies. Google's forays into AI have yielded Nobel-prize winning improvements in medical technology.^(42,43) Both companies have world-class legal teams with decades of experience in litigating intellectual property.

Conclusion

Those of us working in healthcare have long lamented that Epic is like democracy: it's the worst system of operations, except all the others that have been tried. In an era of breathtaking technological transformation, EMRs remain the expensive, harmful steam engines of an era bygone. The risks of radical transformation are great. Recent advances in AI make the promise greater still.

References

1. Wachter RM, Brynjolfsson E. Will Generative Artificial Intelligence Deliver on Its Promise in Health Care? JAMA. 2024 Jan 2;331(1):65–9.
2. The Second Machine Age [Internet]. [cited 2025 Apr 11]. Available from: <https://wnorton.com/books/the-second-machine-age/>
3. McDuff D, Schaekermann M, Tu T, Palepu A, Wang A, Garrison J, et al. Towards accurate differential diagnosis with large language models. Nature. 2025 Apr 9;1–7.
4. Tu T, Schaekermann M, Palepu A, Saab K, Freyberg J, Tanno R, et al. Towards conversational diagnostic artificial intelligence. Nature. 2025 Apr 9;1–9.
5. Tierney AA, Gayre G, Hoberman B, Mattern B, Ballesca M, Wilson Hannay SB, et al. Ambient Artificial Intelligence Scribes: Learnings after 1 Year and over 2.5 Million Uses. Catal Non-Issue Content. 2025 Mar 31;6(2):CAT.25.0040.

6. Garcia P, Ma SP, Shah S, Smith M, Jeong Y, Devon-Sand A, et al. Artificial Intelligence–Generated Draft Replies to Patient Inbox Messages. JAMA Netw Open. 2024 Mar 20;7(3):e243201.
7. Brynjolfsson E. The productivity paradox of information technology. Commun ACM. 1993 Dec 1;36(12):66–77.
8. Reddy A, Scheinker D. The Case For Mathematical Optimization In Health Care: Building A Strong Foundation For Artificial Intelligence. [cited 2025 Apr 11]; Available from: <https://www.healthaffairs.org/doi/10.1377/forefront.20201110.585462/full/>
9. INFORMS. INFORMS. [cited 2025 Apr 11]. Optimization/Mathematical Programming. Available from: <https://www.informs.org/Explore/History-of-O.R.-Excellence/O.R.-Methodologies/Optimization-Mathematical-Programming>
10. Fonner DE, Buck JR, Banker GS. Mathematical Optimization Techniques in Drug Product Design and Process Analysis. J Pharm Sci. 1970 Nov 1;59(11):1587–96.
11. Introduction to Operations Research [Internet]. [cited 2025 Apr 11]. Available from: <https://www.mheducation.com/highered/product/Introduction-to-Operations-Research-Hillier.html>
12. Chiles CR, Dau MT. An analysis of current supply chain best practices in the retail industry with case studies of Wal-Mart and Amazon.com [Internet] [Thesis]. Massachusetts Institute of Technology; 2005 [cited 2025 Apr 11]. Available from: <https://dspace.mit.edu/handle/1721.1/33314>
13. Yu G. Operations Research in the Airline Industry. Springer Science & Business Media; 2012. 492 p.
14. National Football League Scheduling [Internet]. Gurobi Optimization. [cited 2025 Apr 11]. Available from: https://www.gurobi.com/case_studies/national-football-league-scheduling/
15. Zenteno AC, Carnes T, Levi R, Daily BJ, Dunn PF. Systematic OR Block Allocation at a Large

Academic Medical Center: Comprehensive Review on a Data-driven Surgical Scheduling Strategy. *Ann Surg.* 2016 Dec;264(6):973.

16. Fairley M, Scheinker D, Brandeau ML. Improving the efficiency of the operating room environment with an optimization and machine learning model. *Health Care Manag Sci.* 2019 Dec;22(4):756–67.

17. Al Amin M, Baldacci R, Kayvanfar V. A comprehensive review on operating room scheduling and optimization. *Oper Res.* 2024 Dec 6;25(1):3.

18. Shi Y, Mahdian S, Blanchet J, Glynn P, Shin AY, Scheinker D. Surgical scheduling via optimization and machine learning with long-tailed data : *Health care management science*, in press. *Health Care Manag Sci.* 2023 Dec;26(4):692–718.

19. Why Haven Healthcare Failed. *Harvard Business Review* [Internet]. [cited 2025 Apr 11]; Available from: <https://hbr.org/2021/01/why-haven-healthcare-failed>

20. Amazon Pharmacy | Online Prescription [Internet]. [cited 2025 Apr 11]. Available from: <https://pharmacy.amazon.com/>

21. Amazon closes \$3.9B buy of One Medical | *Healthcare Dive* [Internet]. [cited 2025 Apr 11]. Available from: <https://www.healthcaredive.com/news/amazon-closes-39b-buy-of-one-medical/643245/>

22. Oracle Health [Internet]. [cited 2025 Apr 11]. Available from: <https://www.oracle.com/health/>

23. Capoot A. *CNBC*. 2024 [cited 2025 Apr 11]. Oracle announces new AI-powered electronic health record. Available from: <https://www.cnbc.com/2024/10/29/oracle-announces-new-ai-powered-electronic-health-record.html>

24. The Future of Health [Internet]. [cited 2025 Apr 11]. Available from: <https://www.generalcatalyst.com/stories/the-future-of-health>

25. Landi H. General Catalyst's HATCo to buy Summa Health for \$485M [Internet]. 2024 [cited

2025 Apr 11]. Available from:

<https://www.fiercehealthcare.com/health-tech/general-catalysts-hatco-moves-forward-485m-summa-health-deal>

26. Abdalkareem ZA, Amir A, Al-Betar MA, Ekhan P, Hammouri AI. Healthcare scheduling in optimization context: a review. *Health Technol.* 2021 May 1;11(3):445–69.

27. Jafari H, Salmasi N. Maximizing the nurses' preferences in nurse scheduling problem: mathematical modeling and a meta-heuristic algorithm. *J Ind Eng Int.* 2015 Sep 1;11(3):439–58.

28. Svirsko AC, Norman BA, Rausch D, Woodring J. Using Mathematical Modeling to Improve the Emergency Department Nurse-Scheduling Process. *J Emerg Nurs.* 2019 Jul 1;45(4):425–32.

29. Lin SC, Jha AK, Adler-Milstein J. Electronic Health Records Associated With Lower Hospital Mortality After Systems Have Time To Mature. *Health Aff (Millwood).* 2018 Jul;37(7):1128–35.

30. Particle v. Epic Antitrust Lawsuit [Internet]. [cited 2025 Apr 11]. Available from:

<https://www.particlehealth.com/blog/epic-systems-stranglehold-on-u-s-medical-records-harms-patient-care-lawsuit>

31. Janett RS, Yeracaris PP. Electronic Medical Records in the American Health System: challenges and lessons learned. *Ciênc Saúde Coletiva.* 2020 Apr 6;25:1293–304.

32. List of mergers and acquisitions by Amazon. In: Wikipedia [Internet]. 2025 [cited 2025 Apr 11]. Available from:

https://en.wikipedia.org/w/index.php?title=List_of_mergers_and_acquisitions_by_Amazon&oldid=1275556573

33. Li B, Gilbert S. Artificial Intelligence awarded two Nobel Prizes for innovations that will shape the future of medicine. *Npj Digit Med.* 2024 Nov 25;7(1):1–3.