





Research and Applications

A survey of extant organizational and computational setups for deploying predictive models in health systems

Sehj Kashyap ¹, Keith E. Morse ², Birju Patel ¹ and Nigam H. Shah ¹

¹Stanford Center for Biomedical Informatics Research, Stanford University, Stanford, California, USA and ²Division of Pediatric Hospital Medicine, Department of Pediatrics, Stanford University School of Medicine, Stanford, California, USA

Corresponding Author: Sehj Kashyap, MD, MS, Stanford Center for Biomedical Informatics Research, Medical School Office Building, Room x235, 1265 Welch Road, Stanford, CA 94305-5479, USA (sehjk@alumni.stanford.edu)

Received 16 January 2021; Revised 7 June 2021; Editorial Decision 5 July 2021; Accepted 11 July 2021

ABSTRACT

Objective: Artificial intelligence (AI) and machine learning (ML) enabled healthcare is now feasible for many health systems, yet little is known about effective strategies of system architecture and governance mechanisms for implementation. Our objective was to identify the different computational and organizational setups that early-adopter health systems have utilized to integrate AI/ML clinical decision support (AI-CDS) and scrutinize their trade-offs.

Materials and Methods: We conducted structured interviews with health systems with AI deployment experience about their organizational and computational setups for deploying AI-CDS at point of care.

Results: We contacted 34 health systems and interviewed 20 healthcare sites (58% response rate). Twelve (60%) sites used the native electronic health record vendor configuration for model development and deployment, making it the most common shared infrastructure. Nine (45%) sites used alternative computational configurations which varied significantly. Organizational configurations for managing AI-CDS were distinguished by how they identified model needs, built and implemented models, and were separable into 3 major types: Decentralized translation ($n = 10$, 50%), IT Department led ($n = 2$, 10%), and AI in Healthcare (AIHC) Team ($n = 8$, 40%).

Discussion: No singular computational configuration enables all current use cases for AI-CDS. Health systems need to consider their desired applications for AI-CDS and whether investment in extending the off-the-shelf infrastructure is needed. Each organizational setup confers trade-offs for health systems planning strategies to implement AI-CDS.

Conclusion: Health systems will be able to use this framework to understand strengths and weaknesses of alternative organizational and computational setups when designing their strategy for artificial intelligence.

Key words: artificial intelligence, predictive models, clinical decision support, organizational readiness, healthcare delivery, healthcare organizations, computational infrastructure

INTRODUCTION

Artificial intelligence (AI) and machine learning (ML) enabled healthcare is now a possibility for many healthcare systems that have adopted electronic health records (EHRs) and digitized patient

healthcare data.^{1–4} Yet the process of implementing and integrating AI/ML into clinical workflow remains challenging, and the application of AI/ML clinical decision support (abbreviated as AI-CDS) at point of care is scant.^{5–7} In order to initialize and sustain such tools,

healthcare systems must make up-front investments in technological infrastructure and organizational structures to develop, deploy, and manage these technologies.^{8,9} Yet there has been little description about what configuration options exist or analysis of trade-offs between different strategies.¹⁰ An evaluation of steps taken by early-adopter health systems who have deployed AI-CDS into clinical workflow would help guide other health systems that are developing or accelerating their AI-CDS strategies, facilitate collaboration, and spread adoption of AI-CDS at the clinical bedside.

Integrating an AI-CDS entails decisions around system architecture and governance mechanisms. To develop and deploy AI-CDS tools, healthcare systems have to identify use cases, search for vendor solutions or build homegrown AI-CDS algorithms themselves, and coordinate the technological infrastructure to deliver model results at point of care.^{9,11–13} Decisions are required at the organizational level concerning project governance, how diverse stakeholders are incorporated into design and implementation, and long-term maintenance and quality control.^{12,14} Technologically, AI-CDS tools can run in real time or at set time points, use single or multiple EHR data sources for training and inference, and deliver results natively in the EHR or stand-alone applications designed for new workflows.^{12,15,16} Importantly and differently from traditional CDS, these tools can learn from real-world use, change performance over time, and require ongoing monitoring and vigilance.^{17–19} Given that the cost of deploying a single model into clinical workflow can be high,⁸ setting organizational strategy and competence to do so repeatedly and reliably as healthcare systems transition from pilots to production is likely an important factor for successfully capturing the benefits of AI-enabled healthcare.^{20,21}

Case studies of deployments and efforts to summarize common translational pathways and challenges are helpful in identifying success factors associated with implementing AI-CDS into clinical workflow.^{9,10,22–30} Watson et al highlight common cultural, financial, technological, and design barriers faced by academic health systems deploying AI-CDS into clinical workflows and make recommendations to healthcare systems to overcome them.²³ Orenstein et al lay a roadmap of organizational functions and processes to help healthcare systems more effectively use CDS.²⁰ However, to our knowledge, no study has systematically identified the different types of configurations regarding organizational and computational setups that healthcare systems may adopt for integrating AI-CDS into clinical workflow nor scrutinized their trade-offs.

In this study, we examine the organizational and computational setups currently in use to deploy AI-CDS at point of care to inform investments of health systems initializing AI-CDS efforts.

OBJECTIVES

Our primary objective was to identify the different computational and organizational setups that early-adopter health systems have utilized to integrate AI-CDS into clinical workflows. Secondly, we discuss the advantages and disadvantages each configuration confers for a particular healthcare system investing in AI-CDS.

MATERIALS AND METHODS

Interview outreach

We identified health systems with clinical machine learning development efforts by reviewing submissions to 10 healthcare machine learning conferences and symposia. These included the proceedings

of the American Medical Informatics Association Annual Symposia and Summits; Machine Learning in Healthcare Conference; Health Information and Management Systems, Society Machine Learning and Artificial Intelligence Forum; and the Bay Area Medical Informatics Society annual symposium. In addition, we searched Pubmed for clinical machine learning publications in the past 5 years from academic medical centers with Accreditation Council for Graduate Medical Education (ACGME)-accredited clinical informatics fellowships.³¹ A health system met inclusion criteria if it had publicly shared the status of a clinical machine learning model in planning, deployment, or post-deployment in clinical validation. We did not differentiate between an AI-CDS that was static or had mechanisms to change or retrain on data after deployment.

We contacted the chief medical information officers (CMIOs), associate CMIOs, directors of clinical informatics, or authors associated with relevant publications via email. We invited these individuals to discuss their operational and computational setups for AI-CDS. After additional participants within the institution were identified via snowball sampling,³² interviews were held with 1–2 people from a single institution at the same time in group format. Conversations were structured to extract key, implementable information about the computational and operational setups for AI-CDS in clinical workflow at these sites. Conversations were recorded with permission, and notes were taken for accuracy.

Each interview followed the same overall structure: 1) understanding what models were in deployment, 2) discussion of computational setups to train and run those models in production, and 3) eliciting the organization's current method for managing AI-CDS projects. We used a list of pre-prepared questions, which were adjusted over time based on learnings from prior conversations ([Supplementary Appendix A](#) shows the final semistructured interview guide). In cases where we had multiple informants from the same site, the interview was conducted in group format. Any conflicting information was reconciled during the conversation.

Interview synthesis

After reviewing interview notes, 3 practicing clinical informaticists (SK, KM, BP) iteratively synthesized different organizational and computational setups into meaningful types. We inductively created a classification of organizational setups based on how problems were identified for AI-CDS deployment, how models were developed, what workflows were designed, how trust was established among stakeholders, and how the organization implemented and conducted validation after deployment. We actively considered alternative interpretations of the data and stopped evolving the categories once all 3 reviewers agreed on the categories and the classification of each organization. For computational setups, we analyzed what data was used to build the models, how data were sourced at model training and model deployment time, and which computational setup was used to run inference in production. Additionally, we considered how the model output integrated back into clinical workflow. The result was a list of options currently in use for each component. Understanding of an organization's computational setup was informed both by stakeholder interviews and information presented through their publications. Finally, assessment of computational setups specifically intended to facilitate research was beyond the scope of this project.

The Stanford University Institutional Review Board deemed this to be a quality improvement project and exempt from further review.

RESULTS

Thirty-four health systems were identified through publication and conference search. Individuals from all 34 health systems with AI-CDS deployment experience were invited for an interview. Twenty healthcare sites (58% response rate) responded to the invitation, and we interviewed 21 individuals from these sites between February and May, 2020. [Table 1](#) summarizes characteristics of the participating institutions and individuals.

Institutions were geographically located in 11 states. Among individual interviewees, 6 (29%) had roles of Chief/Vice President/Head, 9 (43%) were Directors, 4 (19%) were Professors or Assistant Professors, 1 (5%) was a Clinical Informatics Fellow, and 1 (5%) was a Senior Data Scientist. Seventeen (85%) out of the 19 responders used Epic (Epic Systems, Verona, WI) as their primary EHR vendor, while 2 (11%) used Cerner (Cerner Corp, Kansas City, MO).

The majority of interviewed sites ($n=14$, 70%) had a model running in production that was either homegrown ($n=10$, 50%), sourced from their EHR vendor ($n=12$, 60%), or both ($n=8$, 40%).

There are multiple computational setups used for deploying AI-CDS

Twelve (60%) sites used models developed by their EHR vendor and deployed them through the native EHR vendor configuration, making it the most common shared infrastructure ([Figure 1](#)). Most commonly, models sourced from the EHR vendor were pretrained—did not require the site to use its local patient data for training—used the transactional database as the deployment data source, ran inference in the EHR local or cloud instance, and communicated the model output within the EHR. Computational configurations among the 10 (50%) sites developing and deploying homegrown models varied significantly. Most commonly, the analytic database ($n=9$, 45%), or a research derivative ($n=13$, 65%), were model training data sources while the transactional database was the deployment data source. In addition to these, 4 (20%) sites reported using either their data lakes, non-EHR data sources, and stored streams of transactional data for model training. At model run time, 8 sites (40%) also used their analytic database as a deployment data source and 3 sites (21%) reported using non-EHR and home health data streams. Eight sites (40%) that locally developed models set up a custom data science environment outside their EHR for model inference. Predictive Model Markup Language was commonly used to transfer model weights, particularly for models deployed in local

EHR instances. But this was not the only method, and sites with custom data science environments also wrote custom scripts for data transformation and model inference. Eleven (79%) out of the 14 sites with deployed models visualized model outputs within their EHR, most commonly as automated alerts or columns in patient lists. Five (36%) out of 14 reported to display the outputs outside the EHR, either via secure provider text messaging, web or mobile application, or secure email. Data was exchanged between data sources, data science environment, and communication environment in the form of HL7, FHIR, or EHR-vendor application programming interfaces.

There are 3 organizational configurations used for managing AI-CDS

Organizational configurations for managing AI-CDS were separable into 3 major types: decentralized translation, IT-department led, and AI in healthcare (AIHC). The key distinguishing characteristics between these types was how they identified problems for AI-CDS, sourced, and implemented models. We summarize these details in [Table 2](#).

Decentralized translation

This organizational structure is described as multiple individual research teams working on developing models for clinical use cases. Model needs are typically identified ad hoc. Researchers partner with the health system's IT department to deploy the model, who, in turn, work with researchers to maintain the model. One institution described a faculty research project that developed an AI-CDS to predict sepsis, which was being translated with the help of operations and IT support. We categorized 10 (50%) sites as having a decentralized translation organizational configuration for managing AI-CDS.

IT department led

In IT department led management configurations, the hospital sources models from third-party vendors, or EHR vendors, and deploys models with the help of its IT department. Committees including frontline healthcare workers, subject matter experts, and evaluation specialists are created to oversee the process of implementation and integration. Homegrown models are less commonly used by these sites. Two sites (10%) exhibited this organizational structure.

AI in healthcare team

Sites that use AI in healthcare teams form multidisciplinary groups comprising machine learning engineers, software and database architects, clinicians, and operational leaders. These groups may centralize model development similar to IT department led operations, but unlike IT department led operations, they have the capability to locally develop models as in decentralized translation. Eight (40%) of sites were in the process of developing or had developed AIHC units at their health system.

Table 1. Characteristics of informants and their institutions

Total interviewed sites ($n=20$)	
Accredited Informatics Fellowship	12 (60%)
Ranked in US News Top 20	8 (40%)
Academic medical center	20 (100%)
Seniority of Informants ($n=21$)	
Executive/Senior	17 (81%)
Non-executive/Non-senior	4 (19%)
Sites by number and phase of models ($n=20$)	
One or more in development	20 (100%)
One or more in production	14 (70%)
Source of models in production ($n=14$)	
EHR vendor	12 (60%)
Homegrown	10 (50%)
Both	8 (40%)

DISCUSSION

We describe the current state of computational configurations and organizational setups being used by early-adopter health systems that are integrating AI-CDS into clinical workflow. We find that there is no singular computational configuration that enables all the use cases for AI-CDS. While most surveyed health systems use the default computational configuration provided by their EHR, several extend it by building ancillary components to enable specific use

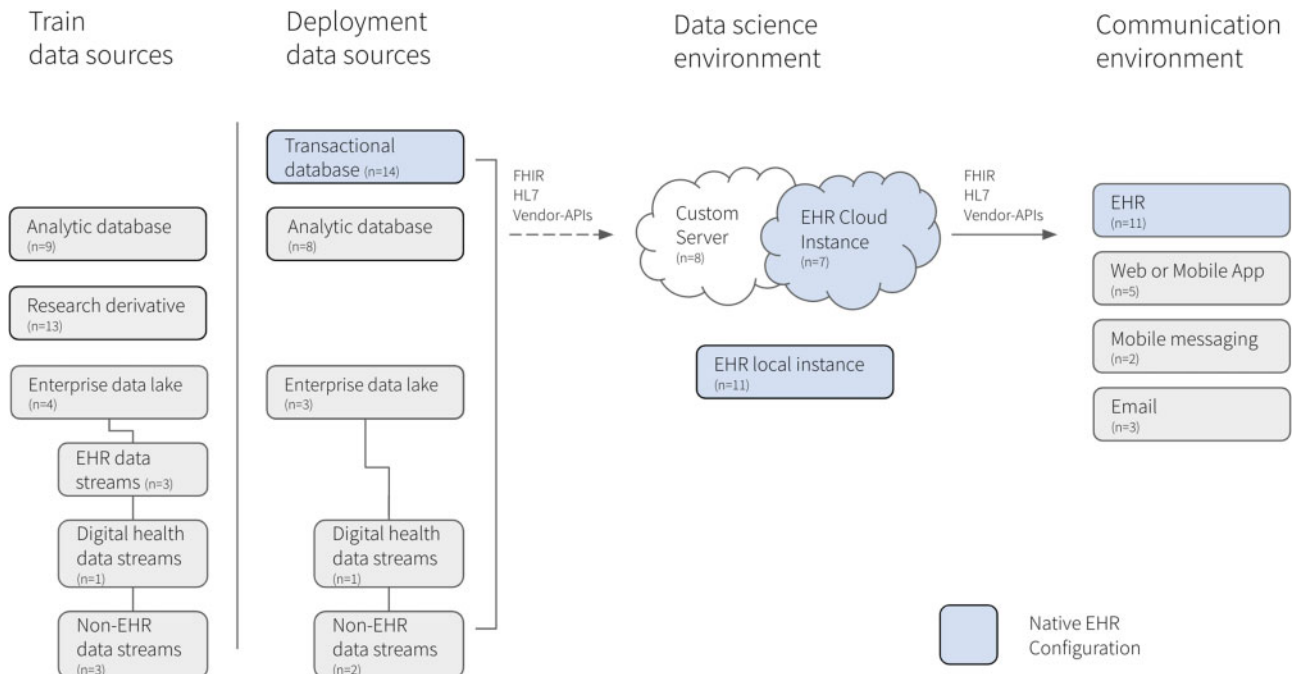


Figure 1. Numerous reported data sources for model training and deployment time and multiple environments for running model inference and displaying results. API: application programming interface; EHR: electronic health record; FHIR: Fast Healthcare Interoperability Resources; HL7: Health Level Seven.

Table 2. Characteristics of 3 organizational structures

	Decentralized translation (N = 10; 50%)	IT Department-led (N = 2; 10%)	AIHC Team (N = 8; 40%)
Definition	Researchers home-grow models and work with IT to integrate into workflow	Hospital IT department sources models from 3rd parties and models are not homegrown	Dedicated team for establishing AI-driven workflows. Can home-grow models or source them from vendors.
Problem Identification	Researcher identified	Committees + outside model availability	Community RFA and ticketing process
Model Builder	Clinicians, ML research groups	External company. Example: EHR vendor	AIHC ML researchers
Model Implementer	IT department	IT Department	AIHC team + IT department

cases. One site developed a system to deliver AI-CDS via mobile messaging to doctors in outpatient centers; another reported displaying it on a web application on a tablet that a triage nurse could carry in the emergency department. Unlike traditional CDS, which is displayed predominantly in the EHR, AI-CDS may need to communicate outside the EHR as it enables new workflows.^{14,33} Off-the-shelf computational infrastructures may represent the easiest entry into deployed AI-CDS; however, health systems need to consider their desired applications for AI-CDS and whether investment in extending the default infrastructure is needed or feasible to accomplish their goals. Additionally, vendors of AI-CDS may need to adjust their architectures to work within a computational setup. In particular, analytic databases were the most common source of data at training and deployment time; this is typically created through nightly extract, load, and transformation processes done on the transactional database. Models trained on analytic databases may perform differently or require significant work on data binding if run off real-time data feeds from the transactional database. Health systems that have use cases for truly real-time AI-CDS need to

closely inspect performance prospectively or reconsider whether their use case requires real-time inference.

In organizational structures, we recognized 3 types of patterns: decentralized translation, IT-department led, and AIHC team driven. Most sites were organized in the decentralized translation pattern. This organizational type enables a high volume of novel AI-CDS to be developed by individual research teams, but implementation is dependent on coordination with IT departments. While most deployed AI-CDS is in a nascent phase, concerns have been raised about AI-CDS longevity due to ongoing maintenance requirements brought on by changing workflows and data drift.^{27,29,34–36} In decentralized translation, there may be a higher probability of project rejection or abandonment if long-term support and maintenance prove difficult to organize. In comparison, health systems adopting an IT department-led pattern centralize sourcing AI-CDS and deployment. Here, the health system tends to deprioritize developing models internally and use third-party model vendors to meet its use cases. This likely increases durability of the AI-CDS, while development of novel models is traded off. The AIHC pattern enables novel

model development and increases the likelihood of long-term maintenance through the model life cycle by a central team managing implementation. Out of the 3 operational structures, dedicated AIHC teams more often extended the native EHR configurations, whereas IT-led operations primarily used native EHR vendor configurations. Even as commercialization and adoption of AI-CDS proliferates, health systems may continue to invest in AIHC teams to develop novel models for the long tail of specific use cases, localize models to their populations, and design new AI-enabled workflows.^{14,37}

Health systems in the decentralized translation operations may be in an exploratory phase of deployment that precedes eventual transition to 1 of the other 2 operational identities. Analogously, early computerized physician order entry and CDS for adverse events and dosage calculation were developed in academic medical centers and diffused through commercial vendors.^{38–40} Traditional non-AI CDS is now primarily developed by EHR vendors, deployed via EHR and managed by IT departments with clinical support.²⁰ If AI-CDS follows the trajectory of traditional CDS, the role of the EHR and IT department will likely grow while researcher roles decrease. Institutional investment will be key to transforming decentralized translation operations into either subtype of AIHC or IT department led. Both AIHC and IT department led organizational structures may represent approaches to achieving organizational readiness to implement AI-CDS.²¹ Health systems must decide if they have resources or need to create AIHC teams or if their use cases are met with existing vended models and the IT department-led approach is preferred.

Our data suggests that there is no universally accepted way to deploy AI-CDS. Early decisions that health systems establish different phenotypes of organizational and computational setup and separate them into different categories. These reflect different priorities and expectations the system has for AI-CDS. Recognizing their own phenotypes can help health systems prioritize investments and build a learning network with the right peer organizations.

Limitations

There are multiple limitations of our study. In particular, our study was not set up to be an exhaustive survey, and there may be operational and computational infrastructures that are not represented here. Most of the institutions that we interviewed were academic medical centers (AMCs). While we did not target these institutions specifically, our list was likely enriched due to our methodology of discovering institutions. These institutions have resources and expertise to develop and deploy models, and their strategies may not represent that of other hospitals. Similar to prior studies, the experiences of non-AMCs are not captured in the current study. Non-AMCs may place less emphasis on home-grown models⁴¹ and have IT-department led organizational configuration and primarily use native EHR computational configuration to deploy models. Additionally, our interviewee selection biases towards health systems that publish on their experience. Given the complexities of large health systems, we may be missing details that could change our characterization of a site's computational and organizational configurations. We are reliant on reported capabilities of health systems, which are difficult to verify. Finally, we do not differentiate between static and self-improving AI-CDS models which learn and change their performance with use, and the latter need to be studied further as they become adopted to see what implications there are for organizational and computational infrastructure.

Further research is necessary to quantify the ultimate return on investment (ROI) of 3 organizational setups, as we are unable to make definitive comparisons here. For example, health systems with AIHC-driven operations certainly publish more about their deployments, but we also see IT-led departments deploying numerous models but not prioritizing academic publishing. The value of this work is identifying options available to health systems investing in AI-CDS and how early decisions in organizational and computational configuration can set organizations onto similar, predictable paths.

CONCLUSION

We characterized computational setups at health systems for deploying machine learning models at the point of care, and the organizational structures that manage the development and deployment of AI-CDS. We identified a diverse range of health systems deploying machine learning models, enabled by multiple computational setups, and managed by 3 types of organizational structures. Institutions striving for demonstrable return on the use of AI-CDS must consider investments in both organizational support and computational infrastructure that align with their priorities.

FUNDING

This work was supported by Stanford Health Care, The Department of Medicine in the Stanford School of Medicine, and the Debra and Mark Leslie endowment for AI in Healthcare.

AUTHOR CONTRIBUTIONS

All authors were involved in drafting or critically revising the presented work, gave final approval of the submitted manuscript, and agree to be accountable for ensuring the integrity of the work. SK, KM, and BP drafted the manuscript, which was reviewed and edited by all coauthors. SK, KM, and NHS envisioned the study. SK and KM interviewed and acquired the data to analyze. SK, KM, and BP conceived of the analysis design and characterized the computational infrastructures and organizational configurations into subtypes.

DATA AVAILABILITY

The data underlying this article cannot be shared publicly to protect the privacy of individuals included in the study.

SUPPLEMENTARY MATERIAL

[Supplementary material](#) is available at the Journal of the American Medical Informatics Association online.

ACKNOWLEDGMENTS

We would like to thank the members of the Stanford Center for Biomedical Informatics Research for reviewing earlier iterations of this work.

CONFLICT OF INTEREST STATEMENT

None declared.

REFERENCES

- Rothman B, Leonard JC, Vigoda MM. Future of electronic health records: implications for decision support. *Mt Sinai J Med* 2012; 79 (6): 757–68.
- Obermeyer Z, Emanuel EJ. Predicting the future - big data, machine learning, and clinical medicine. *N Engl J Med* 2016; 375 (13): 1216–9.
- Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA* 2018; 319 (13): 1317–8.
- Escobar GJ, Liu VX, Schuler A, et al. Automated identification of adults at risk for in-hospital clinical deterioration. *N Engl J Med* 2020; 383 (20): 1951–60.
- Emanuel EJ, Wachter RM. Artificial intelligence in health care: will the value match the hype? *JAMA* 2019; 321 (23): 2281–2.
- KLAS. Healthcare AI: Investment Continues But Results Slower Than Expected; 2020. <https://klasresearch.com/report/healthcare-ai-2020/1443> Accessed December 13 2020.
- U.S. Government Accountability Office. Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care; 2020. <https://www.gao.gov/assets/720/710920.pdf> Accessed December 13, 2020.
- Sendak M, Gao M, Nichols M, et al. Machine learning in health care: a critical appraisal of challenges and opportunities. *EGEMS (Wash DC)* 2019; 7 (1): 1.
- Chen J, Chokshi S, Hegde R, et al. Development, implementation, and evaluation of a personalized machine learning algorithm for clinical decision support: case study with shingles vaccination. *J Med Internet Res* 2020; 22 (4): e16848.
- Drysdale E, Dolatabadi E, Chivers C, et al. Implementing AI in healthcare. *vectorinstitute.ai*. 2020. <https://vectorinstitute.ai/wp-content/uploads/2020/03/implementing-ai-in-healthcare.pdf> Accessed December 1, 2020.
- Sendak MP, Ratliff W, Sarro D, et al. Real-world integration of a sepsis deep learning technology into routine clinical care: implementation study. *JMIR Med Inform* 2020; 8 (7): e15182.
- KLAS. Healthcare AI 2019 Actualizing the Potential of Artificial Intelligence; 2019. <https://klasresearch.com/report/healthcare-ai-2019/1291> Accessed December 13, 2020.
- OPTUM. 3rd Annual Optum Survey on AI in Health Care; 2020. <https://www.optum.com/business/resources/ai-in-healthcare/2020-ai-survey.html> Accessed December 13, 2020.
- Li RC, Asch SM, Shah NH. Developing a delivery science for artificial intelligence in healthcare. *NPJ Digit Med* 2020; 3: 107.
- Sendak M, Elish MC, Gao M, et al. ‘The human body is a black box’ supporting clinical decision-making with deep learning. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*; 2020: 99–109.
- Timsina P, Kia A. Real-Time Machine Learning Pipeline: A Clinical Early Warning Score (EWS) Use-Case; 2019. <https://365.himss.org/sites/himss365/files/365/handouts/552564828/handout-62.pdf> Accessed December 10, 2020.
- Petersen C, Smith J, Freimuth RR, et al. Recommendations for the safe, effective use of adaptive CDS in the US healthcare system: an AMIA position paper. *J Am Med Inform Assoc* 2021; 28 (4): 677–684.
- FDA. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning-Based Software as a Medical Device: Discussion Paper and Request for Feedback. 2019. <https://www.fda.gov/media/122535/download> Accessed December 1, 2020.
- FDA. Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan; 2021. <https://www.fda.gov/media/145022/download> Accessed December 1, 2020.
- Orenstein EW, Murthu N, Weitkamp AO, et al. Towards a maturity model for clinical decision support operations. *Appl Clin Inform* 2019; 10 (5): 810–9.
- Alami H, Lehoux P, Denis J-L, et al. Organizational readiness for artificial intelligence in health care: insights for decision-making and practice. *J Health Organ Manag* 2020; 35 (1): 106–14.
- Parikh RB, Manz C, Chivers C, et al. Derivation and implementation of a machine learning approach to prompt serious illness conversations among outpatients with cancer. *JCO* 2019; 37 (suppl 31): 131.
- Watson J, Hutyra CA, Clancy SM, et al. Overcoming barriers to the adoption and implementation of predictive modeling and machine learning in clinical care: what can we learn from US academic medical centers? *JAMIA Open* 2020; 3 (2): 167–172.
- Sendak MP, D’Arcy J, Kashyap S, et al. A Path for Translation of Machine Learning Products into Healthcare Delivery. *emj.emg-health.com*. <https://www.emjreviews.com/innovations/article/a-path-for-translation-of-machine-learning-products-into-healthcare-delivery/> Accessed August 17, 2020.
- Henry K, Wongvibulsin S, Zhan A, et al. Can septic shock be identified early? Evaluating performance of a Targeted Real-Time Early Warning Score (TREWScore) for septic shock in a community hospital: global and subpopulation performance. In: *D15. Critical Care: Do We Have a Crystal Ball? Predicting Clinical Deterioration and Outcome in Critically Ill Patients*. American Thoracic Society; 2017: A7016. doi: 10.1164/ajrccm-conference.2017.195.1_MeetingAbstracts.A7016
- Sandhu S, Lin AL, Brajer N, et al. Integrating a machine learning system into clinical workflows: qualitative study. *J Med Internet Res* 2020; 22 (11): e22421. <https://www.jmir.org/2020/11/e22421/>
- He J, Baxter SL, Xu J, et al. The practical implementation of artificial intelligence technologies in medicine. *Nat Med* 2019; 25 (1): 30–6.
- Wiens J, Saria S, Sendak M, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat Med* 2019; 25 (9): 1337–40.
- Kelly CJ, Karthikesalingam A, Suleyman M, et al. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med* 2019; 17 (1): 195.
- Paleyas A, Urma R-G, Lawrence ND. Challenges in deploying machine learning, a survey of case studies. arXiv [cs.LG]. 2020. <http://arxiv.org/abs/2011.09926>
- Clinical Informatics Fellowship Programs. <https://www.amia.org/membership/academic-forum/clinical-informatics-fellowships> Accessed January 7, 2021.
- Goodman LA. Snowball sampling. *Ann Math Statist* 1961; 32 (1): 148–70.
- Parikh RB, Manz C, Chivers C, et al. Machine learning approaches to predict 6-month mortality among patients with cancer. *JAMA Netw Open* 2019; 2 (10): e1915997.
- Sculley D, Holt G, Golovin D, et al. Machine learning: the high interest credit card of technical debt. 2014. <https://research.google/pubs/pub43146/> Accessed December 10, 2020.
- Subbaswamy A, Saria S. From development to deployment: dataset shift, causality, and shift-stable models in health AI. *Biostatistics* 2020; 21 (2): 345–52.
- Lenert MC, Matheny ME, Walsh CG. Prognostic models will be victims of their own success, unless... *J Am Med Inform Assoc* 2019; 26 (12): 1645–50.
- Cosgriff CV, Stone DJ, Weissman G, et al. The clinical artificial intelligence department: a prerequisite for success. *BMJ Health Care Inform* 2020; 27 (1): e100183.
- Berner ES. *Clinical Decision Support Systems: Theory and Practice*. New York, NY: Springer; 2007.
- Wong HJ, Legnini MW, Whitmore HH. The diffusion of decision support systems in healthcare: are we there yet? *J Healthc Manag* 2000; 45 (4): 240. https://journals.lww.com/jhmonline/fulltext/2000/07000/the_diffusion_of_decision_support_systems_in.8.aspx?casa_token=n7bkAJSilAAAAA:gdLBGuBMLWYDDogG7GO3tUtPOS5MwMRFZ8UENvP GDA6hHqg1lmTjojHKRt3R2qJImmZ_kAQecTdyCy6g-sMloLtN Accessed December 3, 2020.
- Larsen RA, Evans RS, Burke JP, et al. Improved perioperative antibiotic use and reduced surgical wound infections through use of computer decision analysis. *Infect Control Hosp Epidemiol* 1989; 10 (7): 316–20.
- Evans M. Google strikes deal with hospital chain to develop healthcare algorithms. NEWSPLUS,PRO,WSJ-PRO-WSJ.com. 2021. <https://www.wsj.com/articles/google-strikes-deal-with-hospital-chain-to-develop-healthcare-algorithms-11622030401> Accessed May 26, 2021.