


---

## Brief Communications

# 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?

Joshua Watson,<sup>1</sup> Carolyn A. Hutyra,<sup>2</sup> Shayna M. Clancy,<sup>3</sup> Anisha Chandiramani,<sup>4,5</sup> Armando Bedoya,<sup>5,6</sup> Kumar Ilangovan,<sup>4,5</sup> Nancy Nderitu,<sup>7</sup> and Eric G. Poon <sup>4,5</sup>

<sup>1</sup>Department of Surgery, Duke University School of Medicine, Durham, North Carolina, USA, <sup>2</sup>Department of Orthopedic Surgery, Duke University School of Medicine, Durham, North Carolina, USA, <sup>3</sup>Duke Cancer Institute, Duke University School of Medicine, Durham, North Carolina, USA, <sup>4</sup>Division of General Internal Medicine, Department of Medicine, Duke University School of Medicine, Durham, North Carolina, USA, <sup>5</sup>Duke Health Technology Solutions, Duke University Health System, Durham, North Carolina, USA, <sup>6</sup>Division of Pulmonary, Allergy and Critical Care Medicine, Department of Medicine, Duke University School of Medicine, Durham, North Carolina, USA and <sup>7</sup>Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, North Carolina, USA

Corresponding Author: Eric G. Poon, MD, MPH, Division of General Internal Medicine, Department of Medicine, Duke University School of Medicine, Durham, NC 27705, USA; eric.poon@duke.edu

Received 23 July 2019; Editorial Decision 25 August 2019; Accepted 9 October 2019

### ABSTRACT

There is little known about how academic medical centers (AMCs) in the US develop, implement, and maintain predictive modeling and machine learning (PM and ML) models. We conducted semi-structured interviews with leaders from AMCs to assess their use of PM and ML in clinical care, understand associated challenges, and determine recommended best practices. Each transcribed interview was iteratively coded and reconciled by a minimum of 2 investigators to identify key barriers to and facilitators of PM and ML adoption and implementation in clinical care. Interviews were conducted with 33 individuals from 19 AMCs nationally. AMCs varied greatly in the use of PM and ML within clinical care, from some just beginning to explore their utility to others with multiple models integrated into clinical care. Informants identified 5 key barriers to the adoption and implementation of PM and ML in clinical care: (1) culture and personnel, (2) clinical utility of the PM and ML tool, (3) financing, (4) technology, and (5) data. Recommendation to the informatics community to overcome these barriers included: (1) development of robust evaluation methodologies, (2) partnership with vendors, and (3) development and dissemination of best practices. For institutions developing clinical PM and ML applications, they are advised to: (1) develop appropriate governance, (2) strengthen data access, integrity, and provenance, and (3) adhere to the 5 rights of clinical decision support. This article highlights key challenges of implementing PM and ML in clinical care at AMCs and suggests best practices for development, implementation, and maintenance at these institutions.

**Key words:** machine learning, artificial intelligence, qualitative evaluation, predictive models

---

## INTRODUCTION

With the wide implementation of Electronic Health Records (EHRs) in the United States, health care institutions are accumulating high-quality data that reflect the processes and outcomes of care at a rapid rate.<sup>1-3</sup> These data-rich environments combined with the adoption of machine learning techniques have enabled health care organizations to perform robust analyses of clinical data.<sup>4-9</sup> These developments have come at an opportune time, as there is significant interest in and commitment toward investment in analytics to improve care delivery and bend the cost curve. Predictive modeling and machine learning (PM and ML) techniques have the potential to improve care and decrease costs through a variety of mechanisms, including but not limited to early identification of patients requiring more intensive follow-up through readmission and post-operative complication risk models and automation of diagnostic interpretation previously completed by humans such as identification of diabetic retinopathy.<sup>10-12</sup> In some cases, academic medical centers (AMCs) have developed their own algorithms to refine clinical decision-making.<sup>13</sup> Others have relied on models developed by vendors.

Advances outside of the health care industry have further heightened the excitement about the future use of PM and ML. Our daily life exposes us to countless examples of the power of PM and ML algorithms ranging from near real-time translation software to GPS-guided navigation systems. Within health care, early successes of PM and ML models have energized the community.<sup>6,8,9,14-19</sup> Despite this high level of enthusiasm about PM and ML, little is known about the barriers that health care organizations face when they attempt to leverage the emerging fields of PM and ML to optimize care. We hypothesize that AMCs, with their access to clinical and informatics pioneers and experts, might be in the best position to provide insights to the rest of the community on how to overcome these barriers. Therefore, we interviewed key personnel at AMCs involved in clinical PM and ML projects to identify key challenges and characterize best practices.

## MATERIALS AND METHODS

We completed a qualitative study of PM and ML projects at AMCs in the United States. We conducted semi-structured interviews with experts from AMCs across disciplines, professions, and institutional responsibilities. AMCs offer a unique health care environment as they provide resources and expertise outside of a traditional health care setting, including expertise in predictive analytics and machine learning. Thus, for the purposes of this study, we focused on leaders within those organizations who had knowledge of the current state of PM and ML. This study underwent review and approval by the Duke Medicine Institutional Review Board for Clinical Investigations.

### Interview instrument development

For the purpose of our informants and the research team, we defined PM and ML as:

“any method that utilize data mining and statistical techniques that learn from data to offer predictions on outcomes of interest.”

Based on this definition, our multidisciplinary team of 7 clinical informatics investigators, with guidance from a senior informaticist, developed a discussion guide. After a thorough review of the literature on PM and ML use in clinical care, the research team created

over the course of 4 1-h meetings a discussion guide designed to assess the challenges faced and best practices used during each phase of a model's lifecycle (ie development, implementation, and maintenance). The guide, in its final form ([Supplementary Appendix S1](#)), included an introduction and 6 subsequent sections. The sections included questions about: (1) model development, (2) model implementation, (3) model maintenance, (4) challenges, (5) best practices, and (6) overarching thoughts on PM and ML in the clinical context.

### Sampling framework and informant identification

We recruited via email program directors of ACGME-accredited clinical informatics fellowships (<https://www.amia.org/membership/academic-forum/clinical-informatics-fellowships>) to identify informants at their organization to participate in the study. We focused our recruitment on this subset of AMCs based on the assumption that AMCs with the resources and commitment to train the next generation of informatics leaders are more likely to be at the forefront of PM and ML. We sent an introductory email request for 45-min interviews to 33 different institutions. Individuals who did not respond were sent one reminder email. All recipients of the email were encouraged to identify experts at their organization to participate in the interviews. We ceased recruitment of participants when no new themes or concepts emerged in the interviews.

### Code list development

Interviews were conducted and digitally recorded via the online, web-based conferencing service WebEx (Cisco WebEx, Milpitas, CA, USA). The interviews were subsequently transcribed verbatim by a professional transcriptionist. Content analysis of transcripts was conducted using Computer Assisted Qualitative Data Analysis Software, QSR NVivo (version 12) using the grounded-theory approach.<sup>13</sup> Five transcripts were selected to generate iteratively a preliminary code list within the study team until the team reached consensus on the preliminary list and coding conventions.

Each transcript beyond the initial 5 was reviewed and coded independently by at least 2 researchers, who reconciled any differences in coding until consensus was achieved. Proposals for additional codes or modifications to existing codes were reviewed within the research team, and approved changes were adopted by all transcript reviewers. During team meetings, the group also identified emerging key themes. The coding and theme identification process took place across approximately 10 team meetings.

## RESULTS

We sent invitations to 51 individuals across institutions and conducted 21 interviews with a total of 33 individuals across 19 different institutions. These individuals were geographically located across 13 states. The recruitment rate across institutions was 76% (19 institutions interviewed/25 institutions invited). Thirty-two percent of the institutions participating in the project are among the top 20 best hospitals for 2017–2018 according to US News and World Report rankings. [Table 1](#) summarizes the participating institutions and our informants.

As seen in [Table 1](#), informants spanned the country but were heavily clustered in states with large AMCs. The educational background of informants included MD (58%), PhD (9%), data science (27%), and MD/PhD (6%). Sixty-four percent of informants were executives or held leadership positions within their organization.

**Table 1.** Characteristics of informants and their home institutions

Location of institution by region ( $n = 19$ )	
Northeast	8
South	3
Midwest	3
West	5
Number of models in production ( $n = 19$ )	
<3 models in production	13
>3 models in production	6
Educational background of informants ( $n = 33$ )	
MDs	58%
PhDs	9%
Data Science	27%
MD/PhDs	6%
Seniority of informants ( $n = 33$ )	
Executive/senior role	64%
Non-executive	36%

There was wide variability in the use and sophistication of models. Some informants reported they did not have any models in production, while others reported having multiple models in production. Six institutions had 3 or more models in production, while the majority of institutions had fewer than 3 models in production.

The interviews elucidated key barriers institutions face in applying PM and ML techniques. These barriers can be grouped into 5 themes: (1) culture and personnel, (2) clinical utility of the PM and ML tool, (3) financing, (4) technology, and (5) data. Throughout the interviews, informants provided specific interventions and best practices that helped them successfully overcome these barriers.

### Culture and personnel

There were 3 oft-cited cultural challenges to the clinical implementation of PM and ML technology. The first was difficulty building consensus amongst the myriad stakeholders to set a clear direction for tool development. The institutions with greater success building clinically-implemented models incorporated clinicians and other stakeholders into the full development cycle of the model. The clinicians were key to identifying the clinical need, informing the model's development, and helping determine the optimal places for intervention while minimizing disruption to workflow. Perhaps counterintuitively, one interviewee explained that this collaborative approach was more time-intensive and resource-intensive than the technical challenge of building the model. However, the interviewee explained, it is exactly this collaboration that is one of the biggest determinants of whether a model will be successfully integrated into clinical practice. This interviewee said,

...more of the work is actually going to be focused on the intervention and the program to support that intervention in a sustainable way. The tech part, the IT part is, and the analytics part is getting easier and easier you know. They can spin that up and figure that out over a matter of maybe a week or a few weeks.

The second barrier was attitude toward the PM and ML model. Traditionally, clinicians and scientists were comfortable with performance metrics around routine statistical tests. The performance metrics used to evaluate the myriad of ML tools are less familiar to most clinicians. Also, the tools themselves are sometimes "black box" algorithms that cannot be fully dissected and deconstructed even by experts. This, according to some interviewees, has led to healthy skepticism while sometimes limiting clinical

implementation. Some informants reported working with their EHR vendor to clarify methods used to improve transparency into the inner workings of the model. The third barrier to clinical implementation of these models was managing expectations since hype often distorted clinicians' expectations.

Independent of the cultural barriers to clinical implementation of these models, there were personnel limitations. The demand for people skilled in the creation and maintenance of these models is significantly larger than the number of people available to work the models. Shortages and turnover of personnel with the requisite skill to develop the model as well as maintain the model created barriers to implementation. One interviewee captured this challenge as follows:

When people with institutional knowledge move on to other institutions and their institutional knowledge is particularly targeted on the machine learning models that are in production, that creates a knowledge gap and also a sort of responsibility gap that must be filled by someone if these are to be continued.

### Clinical utility

Several informants suggested that models have to be built in a way that is clinically useful. This begins with clear identification of the question that the team seeks to answer as well as the possible interventions before model development. One of the pitfalls some of the groups experienced was to develop models that lacked clear actionability. For example, one interviewee explained of a model:

They validated that higher scores [are] a high risk of readmission, but what we don't know for sure is whether you can do something about it. Someone who has a gazillion illnesses and is super old, they will always have a high risk. Whether it's a modifiable risk is the true test of how useful it is. . .but what we don't know for sure is whether you can do something about it.

Another factor that was identified as affecting clinical utility is the challenge of configuring alerts; striking the right balance between over-triggering and under-alerting when action is needed has proven challenging. This can have a direct impact on the clinical utility of a model. Nursing alarm fatigue, particularly with some of the most critically-ill patients, is a well-characterized phenomenon.<sup>20</sup> One interviewee explained, "one of the biggest challenges in implementation is figuring out what signals you should send and who to send them to, when and how."

Finally, interviewees expressed concern that clinically implemented models might only improve process metrics without impacting clinical outcomes. As one interviewee explained:

How do you prove that an outcome was improved because someone used the dashboard? It's really hard because most of our business surrounds process improvement. So, just because I improved the process doesn't mean I made your disease process better. We think it does, but really there has yet to be good solid evidence that links that.

Some institutions have begun addressing this concern by developing evaluation procedures for their models that measure the clinical impact.

### Financing

Financial concerns varied across institutions. Several AMCs stated that large portions of their work were entirely unfunded, while other interviewees expressed challenges with model development stalling

or being entirely abandoned due to lack of funding. Concern was expressed regarding costs associated with personnel, technical requirements for integration, and expansion of models clinically. Only a couple institutions—those who stated they had the highest number of models in development—indicated no concerns with funding:

Some of the actual developmental work is funded from grants, some of it is done in a research collaboration under an NDA, and some of it is internally developed and internally funded from a variety of different sources, mostly institutional operational funds.

Smaller institutions and those less experienced with the clinical use of models did indicate concerns with funding resources. One of the solutions to this challenge is to have identifiable project leads and clearly-defined endpoints. Use of vendor-supplied PM and ML models provide another avenue for lowering the financial burden on health care institutions.

### Technology

The current state of technology presents a barrier to the creation of clinically-useful models. The technologies themselves are relatively immature in the health care industry. One solution has been to use vendor-developed models with localization of parameters and/or retraining. However, regardless of the method used, integration of these models into the EHR poses a unique challenge described in this interview:

The issue right now is in the more general context of how do you implement decision support in a commercial EHR? And, how do you do it in a standardized way that a third party can build one tool and deploy it across different EHRs? And right now, none of the vendors have really sophisticated ways of embedding decision support both in triggering that decision support and then running decision support and then taking that feedback and bring[ing] it back into the record to do other things. That process is extremely naïve right now.

The technology platform needed to execute and maintain PM and ML models have not been well developed in health care. As one informant explained:

I think that the key challenges in the implementation phase are integration with what will inevitably be legacy technology in an enterprise setting. Technology moves generally at a slower pace than sort of cutting edge, modern tools, and machine learning in particular is one that leverages the most modern tools. Integrating these modern tools with existing frameworks can pose interesting challenges.

### Data

The variability and incompleteness in health care data quality was a barrier. One interviewee said, “the fidelity of the inputs themselves are quite incomplete” while another emphasized that, “people need to trust that the score is accurate and that really only happens when the data is complete.” Further exacerbating the data challenge is the local customization of workflows and system configurations. As one interviewee pointed out, “[The data at] every EMR is different at every health system”.

Independent of access to data, the unstructured nature of the data has limited the development of good models. One of the interviewees described the challenge and helped explain why it is so important to manage expectations:

Most of our data, like everywhere else, is unstructured. The actual theoretical development of new machine learning techniques [that] could deal with noisy, incomplete biased data has been a challenge. We then partner with them [data scientists] to do some really heavy theoretical machine learning work to help, to actually build the models themselves. That work has been taking place both in deep learning, natural language understanding, neural networks, causal inference, etc. The math there doesn't actually exist, and we've been having to create the math to do it.

Another concern noted was that data exists but does not represent the reality of clinical care. For example, one interviewee stated:

There's always enormous work in a project to identify a candidate data [element] that would be useful and then to validate that it's clinically meaningful and relevant. There [are] many, many times where we get started building a model where we pull out data that looks meaningful and carries a lot of signal, and [when] we talk to the clinicians, that data field is laughed at. They don't use it, or it's just something that they have to enter because of billing reasons, and they are horrified that we would actually use that in a model and there's no getting around that.

Even with an optimal dataset, interviewees expressed challenges in maintaining the dataset ranging from validity to completeness.

## DISCUSSION

### Overview

We found several challenges in developing, implementing, and maintaining ML models in the clinical setting, which helped to inform best practices. There was significant variability in the maturity of ML applications across the informant AMCs. There appeared to be near unanimous consensus that the landscape is evolving rapidly, but there are significant challenges in harnessing PM and ML technologies in health care. As demonstrated by the success of many of the informants, the obstacles are not insurmountable. We believe the broader informatics community must embrace a comprehensive approach to fully capitalize on these rapidly-evolving technologies. Based on discussions with our informants and within the research team, we have developed recommendations for the broader informatics community to accelerate the uptake of effective and reliable models. We also developed targeted recommendations for institutions that strive to develop clinical PM and ML applications internally.

### Recommendations for the informatics community

Clinicians have used a range of clinical risk calculators ranging from tools that help determine the importance of providing prophylaxis against venous thromboembolism to prognosticating about various cancers. These tools have gained wider acceptance over time through usage, validation, and uptake by clinicians. PM and ML models are often more complex and most clinicians are not as facile with the performance metrics applied to these models. For example, while most clinicians are comfortable evaluating traditional diagnostic tests based on their sensitivity, specificity, positive predictive value, and negative predictive value, they are less familiar with concepts relevant to evaluating PM and ML models, such as confusion matrices, gain and lift charts, and area under the receiver operating characteristic curve. It is possible the initial excitement and hype surrounding ML models will help overcome these obstacles, but eventually the medical community will be forced to either responsibly evaluate their performance or abandon their usage in

mainstream clinical care. This will require time, education, and perhaps standards to assist in guiding development.

The community may want to establish best practices for model development, implementation, evaluation, and maintenance to help facilitate standardized processes. For example, these best practices might address how and if to include legacy data, the frequency of re-validation after a model is placed in production, and minimum recommended testing prior to generalizing to an expanded patient population. There is also no clear consensus on how to evaluate the models, indicating a need for the broader community to establish standards for performance evaluation. We would suggest a learning collaborative to share best practices and knowledge.

Professional societies and medical schools/training programs should also educate their members on the basics of PM and ML so clinicians have a basic understanding of what they are and how they are applicable, especially since these models cannot address every problem. They could provide opportunities to discuss the design and deployment of models to encourage collaboration across health systems.

Partnership with vendors are also key to building successful ML models. With the widespread adoption of EHRs across AMCs, the data structures for clinical information have been largely defined by vendors. Successful development of models will be incumbent upon understanding the data structures and underlying architecture of the EHR vendors' data repositories. AMCs must work with vendors to create Application Programming Interfaces that enable more rapid innovation. Without deliberate efforts to improve experience and access to health data, the health care industry is at risk of falling behind even more open industries. Even Goldman Sachs decided to allow open access to a once secretly held trading algorithm, recognizing the benefits of "crowd-sourcing" outweighed the benefits of maintaining secrecy (<https://www.wsj.com/articles/goldmans-trading-floor-is-going-open-source-kind-of-11554285602>).

Finally, the informatics community should advise regulators on how to oversee commercialized models. While there is inherent tension between innovation and regulation, the stakes are particularly high in the case of health care. One misstep could result in loss of life and long-term loss of confidence in this emerging technology. It is also important to define what responsibility the clinicians have when relying on clinical decision support (CDS) that is driven by black-box algorithms. Some of the more complex modeling does not lend itself to traditional deconstruction of errors. It will be important to clearly communicate to patients and clinicians the role of the model in driving clinical care. It is important to distinguish between clinical decisions informed by a model versus decisions driven by a model. There are no clear answers to these types of questions, but it is clear we must start wrestling with them now and encourage our patients to participate so we do not undermine public confidence.

### Recommendations for institutions developing clinical PM and ML applications

Institutions can successfully develop, maintain, implement, and evaluate ML models in the clinical realm. In this section, we will outline some of the success factors.

Overall, institutions should establish a governance structure responsible for the oversight of clinical PM and ML activities. Interestingly, the overwhelming majority of informants did not report formal governance structures, but nearly all of them had informal working groups and committees. The structure of these oversight

bodies matters less than their function. The governance structure should provide a mechanism to resolve conflict and prioritize the diverse array of projects. This can improve accountability and help ensure continued funding streams. We found that the people involved in governance played a key role in knowledge management and often disseminated information about models, had visibility on prior institutional successes and failures, and connected the clinical community with the data science community. These functions are critical to the success and longevity of model development.

Data access, integrity and provenance are key to development of models. These organizational assets can only be developed through close collaboration between clinicians, data scientists, and information technology professionals. The institutions with the greatest success in developing models were thoughtful about how to guarantee data integrity. Aggressive validation of pre-existing clinical data was crucial. Some institutions focused on datasets, such as clinical images, that have already met high standards of procurement and storage. A lot of the models utilizing clinical imaging were more advanced than models relying on clinical data procured through less reproducible means. For example, manually entered EHR data displayed high levels of missingness and inaccuracy. Data collected by regulated devices tended to be more complete and more accurate. As such, all clinical data should not be considered equal. Institutions must begin to evaluate the integrity and lineage of their data and characterize them so moving forward models can take into account the strengths and weaknesses of particular data sources.

PM and ML models should be developed in a similar fashion to any CDS tool. They must address the 5 rights for CDS: right information, right person, right intervention format, right channel, and right time in workflow.<sup>21,22</sup> This approach has been fruitful in developing useful CDS tools. There are 4 best practices that could facilitate adherence to the 5 rights of CDS. First, the model development team should work closely with all of the key clinical stakeholders to ensure multiple perspectives are closely considered before development of the model begins. Second, the team must identify a clear intervention that can be executed by its target audience. Third, there must be a pre-identified metrics that can be used to evaluate the performance of the model, with clinical outcome metrics preferable to process metrics. Finally, the team must commit to constant re-evaluation. This re-evaluation should encompass the actual model performance from a technical perspective as well as the clinical utility to practitioners. A highly accurate model ignored by a clinician is not useful, nor is a poorly-performing model that is routinely intervened upon. Successful implementation of these models requires a multidisciplinary team with a long-term commitment.

### Strengths and limitations

The strengths to this study include the breadth of personnel interviewed who hold academic positions across centers with diverse informatics architecture. Our relatively high response rate is another strength. Our team of interviewers and their varied experiences, combined with an iterative approach to content analysis, helped ensure a flexible approach to interpreting data while minimizing any one individual's biases.

There are limitations to our study. Although the discussion guide and flexible interview process enabled broader discussion, there were several topics that could not be discussed as thoroughly as some interviewers and interviewees would have liked due to time

and topic constraints. For example, some interviewees signaled their interest in describing some of the ethical issues encountered by researchers in this space. Others sought to describe in greater detail the specifics by which they developed their models. These topics would warrant separate interviews and publications to address with the appropriate academic rigor. As our interviews were conducted primarily with executive level leaders, we may have missed some perspectives from data scientists and day-to-day clinical users of these models. Our findings may not be generalizable to non-AMCs. Future work should more heavily involve data scientists, clinical staff, and members from the commercial sector to ensure a more complete understanding of recent advances in the PM and ML space.

## CONCLUSIONS

Our interviews suggest there are significant challenges to adopting PM and ML techniques to create useful clinical tools. Institutions appear to face many of the same barriers. This study synthesizes the knowledge and experience of leaders within AMCs to identify ways to overcome these challenges through best practices. This study lays the foundation to begin more collaboration across institutions as they tackle similar challenges.

Although these technologies are rapidly evolving, their application in the clinical realm remains in its infancy. The broader community must embrace these technologies and drive education, patient engagement, and robust standard development to allow for more systematic collaboration across institutions. Institutions must commit to thoughtful and deliberate development of models for the clinical environment. If we do this, gains in the PM and ML field will undoubtedly help shape future medical advances.

## SUPPLEMENTARY MATERIAL

Supplementary material is available at *Journal of the American Medical Informatics Association* online.

## AUTHOR CONTRIBUTIONS

All authors contributed substantially to conception and design of the project and to data analysis. JW, CAH, SMC, AC, KI, and NN conducted all the interviews and coded all the transcripts. JW and EGP drafted the manuscript and all other authors revised it critically for important intellectual content. All authors gave final approval of the version to be published and agree to be accountable for all aspects of the work.

## CONFLICT OF INTEREST STATEMENT

None declared.

## REFERENCES

1. Blumenthal D. Wiring the health system—origins and provisions of a new federal program. *N Engl J Med* 2011; 365 (24): 2323–9.
2. DesRoches CM, Campbell EG, Rao SR, *et al.* Electronic Health Records in ambulatory care—a national survey of physicians. *N Engl J Med* 2008; 359 (1): 50–60.
3. Jha AK, DesRoches CM, Campbell EG, *et al.* Use of Electronic Health Records in U.S. hospitals. *N Engl J Med* 2009; 360 (16): 1628–38.
4. Liu Z, Hauskrecht M. Clinical time series prediction: Toward a hierarchical dynamical system framework. *Artif Intell Med* 2015; 65 (1): 5–18.
5. Peek N, Combi C, Marin R, Bellazzi R. 2015. Thirty years of artificial intelligence in medicine (AIME) conferences. *Artif Intell Med* 2015; 65 (1): 61–73.
6. Beam AL, Kohane IS. Big data and machine learning in health care. *JAMA* 2018; 319 (13): 1317.
7. Murdoch TB, Detsky AS. The inevitable application of big data to health care. *JAMA* 2013; 309 (13): 1351.
8. Harris A. Path from predictive analytics to improved patient outcomes: a framework to guide use, implementation, and evaluation of accurate surgical predictive models. *Ann Surg* 2017; 265 (3): 461–3.
9. Makridakis S. The forthcoming artificial intelligence (AI) revolution: its impact on society and firms. *Futures* 2017; 90: 46–60.
10. Murff HJ, FitzHenry F, Matheny ME, *et al.* Automated identification of postoperative complications within an electronic medical record using natural language processing. *JAMA* 2011; 306 (8): 848–55.
11. Dilisizian SE, Siegel EL. Artificial intelligence in medicine and cardiac imaging: harnessing big data and advanced computing to provide personalized medical diagnosis and treatment. *Curr Cardiol Rep* 2014; 16 (1): 441.
12. Gulshan V, Peng L, Coram M, *et al.* Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* 2016; 316 (22): 2402.
13. Corbin J, Strauss A. Grounded theory research: Procedures, canons, and evaluative criteria. *Qual Sociol* 1990; 13 (1): 3–21.
14. Melton GB, Hripscak G. Automated detection of adverse events using natural language processing of discharge summaries. *J Am Med Inform Assoc* 2005; 12 (4): 448–57.
15. Campbell R. The five “rights” of clinical decision support. *J AHIMA* 2013; 84 (10): 42–7.
16. Hinton G. Deep learning—a technology with the potential to transform health care. *JAMA* 2018; 320 (11): 1101.
17. Sheikh A, Sood HS, Bates DW. Leveraging health information technology to achieve the “triple aim” of healthcare reform: Table 1. *J Am Med Inform Assoc* 2015; 22 (4): 849–56.
18. Erickson BJ, Korfiatis P, Akkus Z, Kline TL. Machine learning for medical imaging. *RadioGraphics* 2017; 37 (2): 505–15.
19. Wang S, Summers RM. Machine learning and radiology. *Med Image Anal* 2012; 16 (5): 933–51.
20. Winters BD, Cvach MM, Bonafide CP, *et al.* Technological distractions (Part 2): a summary of approaches to manage clinical alarms with intent to reduce alarm fatigue. *Crit Care Med* 2018; 46 (1): 130–7.
21. Campbell R. The five rights of clinical decision support: CDS tools helpful for meeting meaningful use. 2013: 6.
22. Bright TJ, Wong A, Dhurjati R, *et al.* Effect of clinical decision-support systems: a systematic review. *Ann Intern Med* 2012; 157 (1): 29.