Machine Learning in Healthcare: What C-Suite Executives Must Know to Use it Effectively in Their Organizations

By Eric Just,
SVP Product Development

Levi Thatcher
VP Data Science

Tom Lawry
Microsoft

Machine learning (ML) is making significant inroads in healthcare, with common predictive models highlighting patients most at risk, improving clinical decision support at the point of care, automating routine clinical tasks, and improving population health management. On the operational side, ML is making its mark in cost, claims, and staff management, and improving patient flow by eliminating bottlenecks.

Despite industry news that might lead you to believe otherwise, it's important to understand that healthcare is just beginning its ML journey. Most organizations are in the early implementation phase of using ML to improve outcomes. This report is intended as a helpful primer on ML in healthcare: It clarifies the difference between artificial intelligence (AI) and ML, outlines basic and advanced clinical, financial, and operational ML use cases, and offers guidance on the most critical aspects of ML to understand (regardless of where you are on the ML trajectory) how to use it effectively.

Machine Learning and Artificial Intelligence Defined

ML and AI are commonly used interchangeably in healthcare, but there are key differences. AI refers to the ability of artificial systems to gain the intelligence required to perform human-like tasks. ML, often seen as a subset of AI that has the greatest interest and traction in healthcare today, leverages data to make predictions in a variety of realms (clinical, operational, financial, etc.). This report focuses on ML and how organizations can use its predictive and prescriptive capabilities to support their clinical, financial, and operational goals.

From Basic to Advanced: Machine Learning's Use Cases

There are myriad clinical, financial, and operational use cases for ML in healthcare, from more common predictive uses cases ready for deployment today, to more advanced prescriptive use cases becoming available as organizations refine their ML capabilities and the systems of intelligence required to support them.

The first step in building an effective ML model is selecting a use case (e.g., reducing the percentage of heart failure patients readmitted within 30 days to improve outcomes and avoid Medicare penalties). As organizations start the ML journey, this should be simple and align not only with business objectives but also data availability. After determining project goals, clinicians or other subject matter experts then work with data analysts or data scientists to identify variables that truly impact outcomes; several models are built and tested by running the dataset through appropriate algorithms. A final model—which is based on one algorithm and one set of input variables—is chosen and saved when its performance is sufficient for the use case.

It's critical that a subject matter expert (like a clinician) works with an analyst on the workflow optimizations necessary to verify that guidance is delivered in a clear, concise, and actionable way. It's also important for this team to vet ML-recommended interventions. Data analysts and clinicians must ask several questions about how an ML model might impact patient care and the clinical workflow:

- Who will be impacted by this use case? (e.g., patients with heart failure).
- Does the model have safeguards in place to ensure accuracy? (e.g., model predictions will be tested for accuracy after 30 days and checked for safeguards again bias [e.g., accuracy is better for one class of patients over another]).
- Which improvement goal will this use case impact? (e.g., reduce the risk of 30-day readmissions).
- How will this use case impact the improvement goal? (e.g., by ensuring that the patients most likely to be readmitted get care management and a smooth transition of care).
- Whose workflow will this use case be implemented in? (e.g., the care manager's workflow).
- At what point in the workflow will the use case be implemented? (e.g., at the frontline where worklists are prioritized).

Clinical Use Cases

The most common clinical use cases for ML in healthcare today are predictive. For example:

- Predicting the survival odds of patients with glioma, a deadly form of brain cancer.
- Assessing CLABSI risk during an inpatient stay.
- Predicting when patients are at risk of an ED visit or readmission by analyzing home monitoring data.
- Forecasting the likelihood of inpatient readmissions within 30 days.

These clinical use cases are examples of starting point for organizations just beginning their ML journeys. While there's a lot to be gained, clinically, operationally, and financially, from predictive use cases like these, there's an increasing interest in (and need for) more advanced prescriptive use cases, such as accurately matching patients with optimal treatment plans, that have far-reaching impacts on population health and precision medicine.

Other clinical use cases include identifying patients who meet the criteria for participation in clinical trials, assisting clinicians with interpreting images used to diagnose conditions, and risk stratifying patients with chronic diseases. Organizations can use ML to improve patient engagement as well; just as ML can provide information to clinicians on optimal treatments for their patients, it can also provide information to patients on optimal behaviors for managing their conditions.

Operational Use Cases

Organizations can use ML to make systemwide operational improvements, from workflow optimization (data-driven recommendations about which patients to prioritize) to staffing optimization (e.g., properly forecasting inpatient or ED surges). One of the most practical operational use cases for ML is prioritizing worklists, which are ubiquitous in healthcare. If an organization's staff can't easily work through their entire worklist, or if a standard intervention doesn't apply to everyone on that list, ML helps prioritize patients who need the most attention and recommends interventions.

Another operational ML use case tackles appointment no-show rates, which are 10 to 30 percent nationwide and lead to under-utilized equipment, reduced provider productivity, and decreased access to care. A rural health system in New England built an ML model (trained on the system's past patients) to predict which appointments were likely to no-show. The model, which looked at variables such as demographics, past no-shows, day of the week, appointment time, appointment type, etc. beat the predictive efficacy of all models published in the literature and created two paths for increased efficiency: (1) outreach for patients who need it and (2) appointments scheduled based on a patient's likelihood of no-showing. Note that more advanced use cases can arise from a relatively basic ML application. It's notably more difficult, for example, to automatically optimize scheduling—via overbooking slots containing folks likely to no-show—compared to using ML to properly order a manual dialer's queue focused on calls to remind and reschedule.

Financial Use Cases

ML benefits extend beyond the clinical and operational realms. One powerful example of ML's financial use cases comes from a large health system in the Midwest suffering from declining margins because many of its patients weren't paying their bills. The system built an ML model (trained on 500,000 accounts) to assign propensity-to-pay risk scores for 140,000 patients per month. The model provided guidance on which patients needed which interventions (e.g., reminder calls, charity care, payment modifications, and alternative sources of insurance). Thanks to this integration, formerly naïve work queues are finally optimized by leveraging past data from that same health system. This type of financial ML solution benefits both patients and health systems: health systems make money they otherwise wouldn't, and patients pay accounts they otherwise couldn't.

Most health systems are pursing more basic, high-level predictive use cases. In the next few years, more organizations will tackle advanced use cases, such as extracting treatment patterns, as their ML capabilities grow and the industry starts leveraging more sophisticated AI techniques (e.g., vision, speech, etc.).

The Ideal Systems of Intelligence for Effective Machine Learning in Healthcare

As the ML market matures and evolves, healthcare organizations must embrace a systems-of-intelligence mentality, in which disparate systems are integrated. Today, most organizations approach ML with data scientists working in silos. As a result, they run into challenges when it comes to operationalizing ML insights. Investing in the right systems of intelligence is essential for automating ML workflows and operationalizing the movement, calculation, and engineering of the data that feeds effective ML algorithms.

Cloud-Based Data Operating System

A cloud-based data infrastructure, like the Health Catalyst Data Operating System (DOS™) solution, combines the features of data warehousing, clinical data repositories, and health information exchanges in a single, common-sense technology platform. DOS is a critical component of the systems of intelligence for implementing ML in healthcare—it is the analytics backbone that enables organizations to manage information in one place, so they can execute on their clinical, financial, and operational use cases.

The industrywide shift from an on-premises data warehouse to a cloud-based data operating system is critical for being able to accomplish more advanced ML use cases, such as predicting the outcomes of patients with rare diseases. Organizations need a broader dataset from disparate data sources (e.g., EHRs, wearables, genomics, hospital sensor data) that paint more complete pictures of patients and their outcomes—and result in more accurate ML models. A cloud-based infrastructure is more agile than an on-premises data warehouse and can store, manage, and manipulate this data more cost effectively.

Organizations need to have this data available and in a format that can be operationalized and pushed into workflows. Yesterday, this looked like a on-premises data warehouse; today, this looks like a cloud-based data operating system with real-time data feeds and the ability to bring insights and algorithms directly into the workflow.

Clinical Workflow Infrastructure

Organizations need a workflow infrastructure that can easily incorporate insights from a data operating system. This capability is a significant departure from what traditional EHRs can do. Traditional EHRs, the main data collection workflow vehicles in healthcare, are not systems of intelligence; they aren't flexible, they aren't designed around patient care (they are designed around producing financial statements), and they aren't built to leverage analytics. EHRs generally miss important data elements (e.g., genomics, social determinants of health, patient-reported outcomes).

The lifecycle of data—the collection, the analysis, the generation of meaningful insights, and then using those insights to guide patient care—depends on the appropriate workflow infrastructure in place; one that is agile, patient-centric, completes the patient picture, and leaves room for innovation in how organizations bring in analytics throughout the data collection process.

Machine Learning in the Real World

Although healthcare is a relatively new player in the ML game, many organizations have been successfully using ML to optimize workflows. Take [Bon Secours Charity Health System] in New York, for example, which used ML to predict which of its patients were likely to be readmitted within 30 days of discharge. In just ten hours, Bon Secours trained an accurate ML model using data on 54,000 past Bon Secours hospitalizations. Once vetted and deployed, this solution resulted in highly accurate ML-powered critical decision support.

It's important to keep in mind that ML models trained with system-specific clinical data are more effective than rules-based models that are based on populations different from those where the tool is used (as patient characteristics and treatment protocols vary by location). For example, the LACE Index, a rules-based model used to predict readmission risk, was developed using data from 4,800 patients (mostly middle-aged patients without serious comorbidities) admitted to 11 hospitals in Canada between 2002 and 2006. Certainly, patients under an organization's care today may differ from those studied in Canada more than a decade ago; fortunately, ML-based models are able to leverage local idiosyncrasies around demographics and process.

Mission Health, a seven-hospital system and ACO in North Carolina, set out to improve on the LACE Index by developing its own predictive model based on its patients, who were significantly different from those in the LACE study. Its model was designed to use data available upon admission (rather than waiting for discharge data, as LACE did). This flexibility around workflow timing is another general benefit of machine learning. Subsequent studies post-deployment showed that Mission's model predicted which patients would be readmitted more accurately than LACE. Mission used the readmission risk predictor to support ongoing improvement in discharge follow-up, reducing its all-cause readmission rate below the level of its peers.

Compliant, Socially Responsible Machine Learning Is Critical

Every new technological frontier has its limitations, and ML is no exception. Healthcare leaders must be aware of these limitations to minimize risk—to their patients and themselves. As organizations start using ML, they need to keep compliance and social responsibility in mind.

Compliance

Most healthcare organizations have data governance committees responsible for prioritizing and monitoring all data-related projects; these groups must closely review, approve, and monitor all ML efforts (as they would with any initiative involving data), including their compliance with HIPPA and alignment with compliance standards. Compliance is critical, but organizations should not let policies and regulations become barriers for pursuing ML in healthcare. Policies rarely encompass all possible use cases, especially with an emerging capability like ML, so organizations must boldly navigate ML regulations and sensibly interpret them as they prioritize bringing ML-powered value to healthcare.

Social Responsibility

Organizations have a social responsibility to guard against ML biases that underserve segments of their populations. For example, an ML algorithm that makes a clinical prediction (i.e., a recommended treatment plan for patients) that benefits one segment of a system's population more than another segment should be mitigated as much as possible. There will always be variables about patients that aren't included in the ML model, but organizations can prioritize socially responsible ML and avoid building historical inequities into their models by relying on clinical expertise and evidence-based research.

Machine Learning in Healthcare: The Accessible Industrywide Imperative

ML can undoubtedly help healthcare organizations and patients by predicting outcomes and recommending interventions to improve those outcomes. ML can also help health systems make financial and operational improvements. While some ML use cases are sophisticated, requiring data from multiple sources, smaller hospitals and ACOs can still use ML to tackle more common use cases, from preventing readmissions to automating routine clinical tasks. These smaller organizations can leverage ML to improve outcomes while continuing to invest in the systems of intelligence required to support more advanced use cases down the road.

ML's promise in healthcare is undeniable, but all things considered, it's is still in its infancy. It's important to understand the most practical use cases for ML today while keeping advanced use cases in sight. It's also important to know that using ML effectively requires robust systems of intelligence that call for significant infrastructure investment, and an awareness of what it takes to create compliant, socially responsible ML. Healthcare leaders with this level of understanding and awareness are positioned to not only take advantage of the basic use cases, but also pave the road for advanced ML use cases in the future.

ML Is a Tool for the Clinician's Decision-Making Process

Finally, it's also important to note that ML is a tool meant to augment (rather than replace) the judgement of clinicians. ML predictions and suggestions for interventions should be treated as another input (albeit an advanced and complex input) for clinician decision making. The goal is that together, ML and clinicians can improve outcomes for patients in an efficient, effective, and sustainable way.