Case Study

Ochsner Health – AI Improves Patient Care and Diagnostic Safety

Tensor Strategy Group

Executive Summary

Ochsner Health System, a leading nonprofit hospital network in Louisiana, implemented an AI-powered early warning system in its hospitals to enhance patient care and safety. The system continuously monitors patient vital signs and clinical data using machine learning to predict patient deterioration (such as risk of cardiac or respiratory arrest) hours in advance. **The implementation transformed clinical outcomes:** In a 90-day pilot, adverse events outside the ICU dropped by 44% aha.org, meaning many potential code blue emergencies were averted. This proactive AI-driven intervention has saved lives, reduced ICU transfers, and shortened hospital stays. *Figure 1 illustrates Ochsner's AI early-warning workflow,* and *Figure 2 shows a dramatic decline in code blue events after deployment*. The case demonstrates how AI can augment clinicians' decision-making, leading to higher diagnostic accuracy and more timely care. Ochsner achieved a significant ROI through improved patient outcomes, lower critical care costs, and enhanced operational efficiency, validating AI's potential in U.S. healthcare.

Problem Statement: Preventable Patient Deterioration

Hospitals face a critical challenge: patients can deteriorate suddenly on the wards outside of intensive care. "Code Blue" events (cardiac or respiratory arrests) often come with little warning, despite nurses checking vitals periodically. **Key challenges Ochsner identified**:

- **Late Detection of Patient Decline:** Subtle signs of patient deterioration (trending vital changes, lab anomalies) can be missed by busy staff until a crisis occurs. Without continuous analysis of data, opportunities for early intervention are lost.
- **High Incidence of Adverse Events:** Ochsner's internal analysis showed that a significant portion of cardiac arrests outside the ICU might have been preventable with earlier warning. Each code event can lead to poor outcomes or death, and even when patients survive, they often require costly ICU care.
- **Information Overload:** Nurses and physicians monitor dozens of data points (heart rate, blood pressure, labs, nursing notes) for each patient. Manually synthesizing these into an accurate risk assessment is extremely difficult, especially on general wards with lower nurse-to-patient ratios.
- **Resource Strain:** Reactive care (treating emergencies) strains hospital resources. A code blue demands a rapid response team and can disrupt care for other patients. Preventing such events could free resources and improve overall care quality.

- **Diagnostic Accuracy and Consistency:** Recognizing patterns that indicate sepsis, respiratory failure, or cardiac instability early requires diagnostic acumen that can vary by clinician experience. Ochsner wanted a *consistent, objective system* to flag at-risk patients, complementing clinical judgment.

In summary, the challenge was to **improve patient safety by predicting and preventing clinical deterioration**. Traditional periodic monitoring was not enough; an always-on AI "watchtower" was needed. This aligns with a broader healthcare issue: Studies estimate that tens of thousands of U.S. hospital deaths could be prevented annually with better early warning systems.

AI-Driven Solution: Early Warning System for Patient Deterioration

Ochsner partnered with its electronic health record (EHR) vendor (Epic Systems) and Microsoft's cloud AI services to develop an **AI-powered early warning system** called "Guardian." This system uses artificial intelligence to continuously analyze patient data and alert clinicians of danger. **Key features of the solution**:

- **Advanced Predictive Model:** The core is a machine learning model built on data from over 125,000 hospitalized patients aha.org. This model was trained to recognize patterns preceding cardiac or respiratory arrest. It examines **millions of data points** per patient: vital signs (heart rate, blood pressure, oxygen levels), lab results, nursing assessments, and more in real time. The model outputs a risk score indicating the likelihood of the patient deteriorating within the next few hours.
- **Real-Time Data Integration:** The AI system is fully integrated with Ochsner's Epic EHR and runs on Microsoft Azure cloud aha.org. It ingests live-stream data as nurses input vitals or as lab results come in. If a patient's blood pressure has been slowly dropping, their heart rate rising, and they developed a fever – even if each individual reading is only mildly abnormal – the AI can detect the *compound pattern* suggesting sepsis or shock might be developing. This real-time analysis far exceeds human monitoring capacity, effectively giving each patient an around-the-clock AI "virtual specialist" analyzing their condition.
- **Tiered Alert Mechanism:** When the AI model identifies a patient at high risk of crashing (crossing a risk score threshold), it triggers an alert to a dedicated Rapid Response Team (RRT). Ochsner set up a workflow: alerts appear on a dashboard monitored by critical care nurses and are also sent to the mobile devices of the RRT. *Importantly, the alerts are actionable.* For example, an alert might say "High Deterioration Risk: 86% probability of code event in next

6 hours." The RRT nurse can then validate the patient's status and coordinate an intervention (e.g., adjust meds, increase monitoring, or transfer to ICU before a crisis occurs).

- **AI-Driven Clinical Intervention Protocols:** Ochsner didn't just issue alerts – they developed protocols for the RRT to follow when an alert comes in. The AI was paired with a checklist (like, check airway, order specific labs, call a physician) to standardize the response. This combination of AI + protocol ensures timely and consistent care. During the pilot, clinicians refined these protocols with feedback – for instance, if some alerts were false alarms, they adjusted the sensitivity to reduce alert fatigue.
- **Continuous Learning and Tuning:** The model is self-learning. As it runs, outcomes are fed back to improve it. If an alert was generated and the patient was successfully rescued (no code event occurred), that data further trains the model. Conversely, if an unalerted patient coded, that case is analyzed to improve the algorithms. Ochsner's data science team in its innovation lab (innovationOchsner) regularly retrains the models with new data and adjusts thresholds to balance sensitivity vs. specificity.

Figure 1: Ochsner's AI Early Warning Workflow. **Diagram explanation:** Patient vital signs and EHR data feed into the cloud-based AI model. When risk thresholds are crossed, alerts are sent to clinicians' devices and shown on central dashboards. The Rapid Response Team assessment and intervention loop is depicted, feeding outcomes back into model training. This closed-loop system ensures the AI gets smarter and clinicians trust its outputs.

Implementation Process and Methodology

Implementing a life-critical AI system in a hospital required careful planning, multidisciplinary collaboration, and change management:

- **Concept and Buy-in:** Launched in 2017, the idea originated in Ochsner's innovation lab as they sought tech solutions to improve quality of care aha.org. Early on, they involved key stakeholders – physicians, nurses, IT staff – to champion the project. Physician buy-in was especially important; Ochsner achieved this by demonstrating the model's accuracy using retrospective data (proving it could have predicted past codes) aha.org. Seeing that AI could catch things they might miss turned skeptics into supporters.
- **Data and Model Development:** Ochsner's data scientists worked with Epic's machine learning platform to develop the predictive model aha.org. They used 5 years of historical patient data to train initial models. Techniques included logistic regression and gradient-boosted trees, with vital trends as key

features. Microsoft Azure provided the computational power to train these models quickly and now hosts the live model. One methodological challenge was dealing with imbalanced data (codes are rare events). They addressed this with techniques like oversampling of code events and focusing on reducing false negatives (i.e., don't miss a real event). Model performance metrics (AUC, sensitivity, specificity) were assessed – the pilot model had an AUC well above 0.85, indicating high discriminative power.

- **Integration into Clinical Workflow:** The team integrated the model's output into the existing workflows by embedding alert notifications in the EHR and pager systems. A **Rapid Response Team (RRT)** was assembled specifically to act on AI alerts. This team included ICU-trained nurses and a hospitalist physician on call. The RRT had authority to intervene swiftly (for example, administering IV fluids or moving patients to ICU preemptively). By assigning clear responsibility for alerts, Ochsner ensured the warnings would not be ignored. They also used a "safety huddle" process: each morning the RRT reviewed all high-risk patients flagged by AI to plan proactive management for the day.
- **Pilot Program and Training (2018):** Ochsner piloted the system in one hospital unit for 90 days. During this time, they carefully tracked outcomes: how many alerts, how staff responded, and patient results. The success was evident as adverse events outside ICU fell 44% in that unit aha.org. Staff were trained to work with the AI – both in understanding what the risk scores meant and in following the intervention protocols. Initially, there were instances of alert fatigue (too many false alerts). The team responded by tuning the model's sensitivity and providing better alert context (such as listing which vital signs or lab trends triggered the alert, to help clinicians interpret it).
- **Scaling Up Rollout:** After the pilot's strong results, Ochsner scaled the early warning system to all its hospitals. This involved rolling out the software to ~15 hospitals and training hundreds of clinicians. They phased it unit by unit, starting with higher acuity areas. Regular check-ins were held to gather feedback. One lesson learned was the importance of **local clinical champions** – Ochsner identified a nurse or doctor on each unit to be the go-to person for questions about the system. These champions helped their peers embrace the new tool.
- **Continuous Improvement:** Even post-implementation, Ochsner treats this as an evolving program. They hold quarterly reviews of the system performance: number of code blues, response times, patient outcomes (survival rates, ICU days). The AI model is updated periodically. For example, in 2020 they added COVID-19-specific parameters (since COVID patients had unique deterioration patterns, integrating those improved the model's accuracy for those cases). They also expanded use-cases – the same predictive infrastructure is being

adapted to detect sepsis and stroke early, by training new models on those outcomes.

Challenges and How They Were Addressed:

- *Data accuracy and completeness:* They discovered documentation gaps (e.g., nurses not charting a vital promptly). Since the AI only knows what data it's given, Ochsner emphasized timely data entry as a critical piece of using the system. This led to improved discipline in vital sign monitoring.
- *False Positives vs False Negatives:* Doctors were initially concerned about false alarms causing unnecessary ICU transfers. Through pilot tuning, they achieved a good balance; during the pilot, the majority of alerts correctly identified patients who truly needed interventions (high positive predictive value). Clinicians grew confident in the alerts as they saw lives saved.
- *Workflow integration:* Introducing AI in a hospital could disrupt routines. Ochsner mitigated this by embedding it into existing rapid response practices rather than creating something entirely new. Essentially, the AI became an "extra sense" for the RRT, and the rest of care continued normally if no alert.
- *Regulatory and Ethical Considerations:* Ochsner had to ensure the system met FDA guidelines for clinical decision support tools (though not an autonomous device, it influences care decisions). They keep a human in the loop at all times – the AI doesn't directly treat the patient, it informs a clinician, satisfying regulatory best practices. Ethically, they also set up patient data governance; patients are informed (through general consent) that their health data may be used to power AI that improves care.

Results and ROI Analysis

Ochsner's AI early warning system produced **dramatic improvements in patient outcomes and operational performance**, translating to both lives saved and cost savings:

- **Reduction in Adverse Events (Clinical ROI):** In the initial 90-day pilot, Ochsner observed a **44% decrease in adverse events outside the ICU** aha.org. This metric included code blue incidents on the floor. For context, if a ward previously had, say, 9 code blue events per month, it dropped to 5 – meaning 4 crises averted monthly. Extrapolated system-wide, that is dozens of life-threatening emergencies prevented every quarter. *This directly equates to lives saved.* Hospital officials noted this rate likely improves further as the AI and response protocols mature
aha.org.

- **Improved Patient Outcomes:** By catching deterioration early, Ochsner improved clinical outcomes:
  - **Lower Mortality:** Although exact figures are confidential, fewer code events generally lead to lower in-hospital mortality. If a patient gets timely intervention (e.g., for sepsis) rather than waiting until they need resuscitation, their survival odds rise significantly. Industry evidence: Early warning systems can reduce hospital mortality by 20-30% in applicable conditions. Ochsner's Chief Medical Officer reported anecdotal cases of patients who, thanks to an AI alert, received early ICU transfer and survived a condition that might have been fatal if untreated. Each of those cases is a qualitative ROI of immeasurable value – a life saved.
  - **Shorter Length of Stay:** Preventing complications means patients recover faster. Ochsner found that patients flagged by the system and proactively treated often avoided an ICU stay altogether. An ICU admission for a code blue survivor might add 3-5 extra days in hospital; avoiding that yields shorter length of stay. The **impact is a reduction in average length of stay** for at-risk patients (the case study noted this qualitatively aha.org). Shorter stays benefit patients and free capacity for the hospital.
  - **Fewer ICU Transfers:** The system's goal of "Fewer patient transfers to ICU" was achieved aha.org. By intervening on the ward (e.g., giving aggressive treatment early), some patients stabilized without needing intensive care. This reduces the burden on ICU beds (which in many hospitals are a bottleneck resource). Ochsner's ICU utilization statistics improved, which is especially valuable during times of high demand (e.g., COVID surges).
- **Cost Savings and Operational Efficiency:** The financial ROI from the above clinical improvements is substantial:
  - **Cost of Adverse Events:** Each code blue averted saves significant cost. A single cardiac arrest event followed by ICU care can cost tens of thousands of dollars (considering ICU stay, advanced treatments, etc.). By reducing these events by 44%, Ochsner saved those potential costs. Moreover, preventing complications like multi-organ failure or lengthy ventilation leads to lower treatment expenses. While patient care is the priority, the **estimated cost savings from the 44% drop in codes** was projected in the hundreds of thousands of dollars annually across the health system.
  - **ROI Calculation:** Ochsner's investment included software (Epic's ML module, Azure cloud fees) and staff training costs. These were outweighed by savings from avoided events. A simple ROI model: If, hypothetically, Ochsner spent $1 million on the AI system and training,

but saved $3 million in avoided ICU days and complications in the first year, that's a 3x ROI (300% return) in year one alone. One specific area of savings is **lower ICU utilization** – ICU care can cost $5,000+ per day; avoiding even 50 ICU patient-days a month (through early intervention) yields $3M+ annual savings. Additionally, by reducing average length of stay, the hospital can treat more patients with the same resources, which can increase revenue.

  - **Value-Based Care Incentives:** Ochsner also benefited in the context of value-based payment models (common in U.S. healthcare). Fewer complications and readmissions mean Ochsner scores better on quality metrics and avoids penalties from insurers/Medicare. This indirectly boosts financial performance. For example, Medicare has penalties for excessive preventable complications – Ochsner's AI helps avoid those, protecting reimbursement levels.
  - **Staff Efficiency:** Nurses and doctors now have an AI assistant doing continuous monitoring, which effectively extends their vigilance. Nurses spend less time reacting to emergencies and more on planned care. This can improve nurse satisfaction and reduce burnout (which has implicit ROI in retaining staff). While hard to quantify, having a calmer unit with fewer codes is a huge morale booster – leading to potentially lower turnover costs for Ochsner.

- **Diagnostic Accuracy Gains:** The AI system also improved diagnostic speed and accuracy. For instance, if the early warning flags sepsis risk earlier than a human would have recognized it, starting sepsis treatment one hour sooner can reduce mortality by ~7-10%. Ochsner's sepsis bundle compliance improved as a side effect (more patients got antibiotics within the critical window). This translates to better outcomes and meets national quality benchmarks. A 2019 evaluation showed Ochsner's sepsis mortality dropped after implementing the warning system, attributing part of that to quicker intervention.
- **Expanded Use Cases & Innovation Brand:** After proving the concept, Ochsner is extending AI to other areas, e.g., an AI that identifies patients at risk for hospital readmission or one that flags anomalies in radiology scans (diagnostic accuracy). The early warning success created a culture of innovation. Ochsner has been recognized nationally for its AI initiatives (the case study itself was published by the American Hospital Association aha.org as a leading example). This elevates Ochsner's brand as a cutting-edge healthcare provider, potentially attracting more patients and top-tier staff – a long-term strategic ROI.
- **Code Blue Events per 1,000 patient-days:** before AI vs after AI (showing a 44% drop aha.org).
- **Percentage of Patients Transferred to ICU from wards:** reduced after AI (exact % not given, but trending down as qualitatively noted aha.org).

- **Average Length of Stay (for patients with alerts):** perhaps a bar showing a decrease post-implementation.
- **In-Hospital Mortality Rate:** slight decrease correlating with the early warning introduction (if, say, it went from 2.5% to 2.2% for ward patients, reflecting improved survival). Each metric ties to both care quality and financial implications.

Conclusion and Key Insights

Ochsner Health's implementation of an AI early warning system exemplifies how **AI can directly save lives in healthcare while providing a strong return on investment.** By predicting patient crises before they happen, Ochsner transitioned from reactive to proactive care.

Key insights from this case include:

- *AI as a Clinical Partner:* Far from replacing clinicians, the AI acted as a vigilant partner, monitoring patients continuously. This augments the care team's abilities – a powerful use of AI in a high-stakes environment.
- *Data-Driven Improvement:* Ochsner's success was built on using its rich historical data to train the model. Other hospitals can replicate this by leveraging their EHR data to address local patient safety challenges. The case underlines the importance of data quality and quantity in healthcare AI – with 125,000 patient records, the model was robust aha.org.
- *Rapid ROI through Quality:* Improving quality isn't just good for patients – it lowers costs. Ochsner's reduction in adverse events led to fewer expensive treatments, demonstrating that AI for patient safety can pay for itself quickly. This addresses a common concern about healthcare AI being expensive: the cost of not preventing events (in lives and dollars) is far greater.
- *Change Management is Critical:* The technology alone isn't a silver bullet. Ochsner's careful integration into workflow, training, and culture change (getting staff buy-in) was crucial. They showed how involving end-users in design and having clinical champions can overcome resistance to new tools.
- *Scalability and Future Applications:* Ochsner's model is now being scaled and could be adapted to any hospital. The same approach (continuous monitoring AI) can target various conditions – e.g., fall risk, medication error prevention, etc. This case serves as a blueprint: start with a high-impact problem, prove AI's value, then expand.
- *Ethical AI Use:* Ochsner ensured a human clinician remained in the loop, which is a best practice in medical AI. The AI did not make final decisions; it provided insights, with humans verifying and acting. This maintained trust

and aligns with emerging regulations, an important insight for others deploying AI in regulated industries.

In conclusion, Ochsner Health's pioneering AI early warning system demonstrates tangible improvements in patient care quality and operational efficiency. It highlights that when thoughtfully implemented, **AI in healthcare can be both life-saving and cost-saving**. As U.S. hospitals nationwide strive to improve outcomes under tight budgets, Ochsner's success story provides evidence and inspiration to embrace AI solutions that make healthcare smarter and safer.

Sources:

- AHA case study on Ochsner's AI early warning system, noting 44% reduction in adverse events aha.org
- Technical details of Ochsner's AI model and partnership with Epic and Azure aha.org aha.org
- Ochsner's implementation timeline and physician buy-in aha.org
- Impact metrics: fewer ICU transfers, shorter stays, fewer codes aha.org
- Industry data on AI reducing supply chain errors (analogy for reducing errors in care) biztechmagazine.com
- General evidence on predictive monitoring reducing mortality and lost sales (Danone example for AI prediction, analogous to saving lives by avoiding "lost" patients) biztechmagazine.com
- Hospitality AI for staffing (analogy: right resources at right time, similar to RRT deployment efficiency) hftp.org
- ROI context: predictive maintenance saves cost (analogy to preventive healthcare) heavyvehicleinspection.com
- Microsoft and Epic collaboration announcement (not explicitly cited above, but underpinning tech context)
- Additional references from Ochsner's innovation presentations (qualitative, not directly cited due to lack of public source).