Implementing AI for Patient Privacy at IU Health

Mitchell Parker, Executive Director, Information Security and Compliance, IU Health

HIMSS
CENTRAL & SOUTHERN OHIO *Chapter*

Purpose of Presentation

To discuss the realistic steps needed to implement an effective proactive patient privacy monitoring and response system

Agenda

- Why did we do this?
- What incorrect assumptions do people make?
- Risk Assessment
- How the journey began – starting with HR
- Operational Staffing Plan
- Policy Changes – what policies do you need?
- Communication Plan for Leadership
- Organizational Communication Plan
- Implementation and Monitoring
- Continual Follow-up and Metrics
- How this influenced further AI/ML/Deep Learning work for us?

Why did we do this?

- We have multiple large electronic medical record systems for storing patient data
  - Epic, Cerner Millennium, Meditech, PACS, Allscripts, etc.
  - Not just the EMR, but also Population Health and Patient Portals now!
- We have requirements under the HIPAA Security Rule and Meaningful Use Program to inspect audit log files to ensure no unauthorized accesses of patient data occurred
- Unauthorized accesses are considered data breaches under the HIPAA Privacy and Security rules

Why did we do this?

- EMR and EHR systems generate millions of access logs a day
  - A large health system such as ours can generate 15 million on just the EMR alone
  - There is no way you are going to have manual review or statistical sampling catch these events
- For Riley Children's Health, which sees 300,000 patients a year, to get to a 95% confidence level with a 5% confidence interval , you'd have to look at 378 records a month manually
  - Some of these records can have tens of thousands of entries
  - Dr. Bimal Desai, CHIO at CHOP, cited that they can generate almost 100,000 entries for a single child's two week stay (https://medcitynews.com/2014/07/childrens-hospital-cmio-turns-entrepreneur-to-prevent-patient-data-breaches/)
  - There's no way you're going to look for anything meaningful there – even with hundreds of pairs of eyes

Why did we do this?

- The HIPAA Security Rule has 3 places that require it:
  - §164.308(a)(5)(ii)(C) Addressable Does your practice include log-in monitoring as part of its awareness and training programs?
  - §164.312(b) Standard Does your practice have audit control mechanisms that can monitor, record and/or examine information system activity?
  - §164.308(a)(1)(ii)(D) Required Does your practice regularly review information system activity?

What did we pick?

- After looking at several solutions, we picked Haystack Informatics' product
  - Since acquired by Iatric Systems
  - We also looked at Maize and Protenus
  - We wanted a cloud-based system that was capable of scaling to large amounts of data
    - Haystack uses Apache Redis and sits on Amazon AWS
  - We did not want traditional Relational Database technology as it does not scale to handle billions of records
  - We also had our Red Team pentest the system

Why did we do this?

- We don’t have resources to dedicate to continual monitoring
  - We operate on low margins and low budgets, esp. Privacy!
  - Yet we need to show compliance with the rules
- Security has successfully used automation and machine learning to filter through the hundreds of millions of events and behaviors daily
  - Darktrace - $1B+ British network security company that has a growing municipal presence, including Las Vegas that uses ML to detect threats
  - Blackberry Cylance – Endpoint protection that uses ML to stop unknown attacks
  - IBM QRadar – uses ML to search through large numbers of events
  - JASK – uses AI to search through network and stored events

What incorrect assumptions do people make?

- OCR Fines
  - There's a lot of vendors out there that try and scare people using those as a cudgel
  - Patient Privacy Monitoring is one of many areas they will look at
  - Don't assume that buying this will solve your problems on its own
- You can just plug this in and have it work
  - This is not the case
  - There is a lot of configuration work that has to be done first
  - This is significantly more complex than putting in a security tool
  - 10% of the work is with technology

Incorrect Assumptions

- What people think about AI/ML:

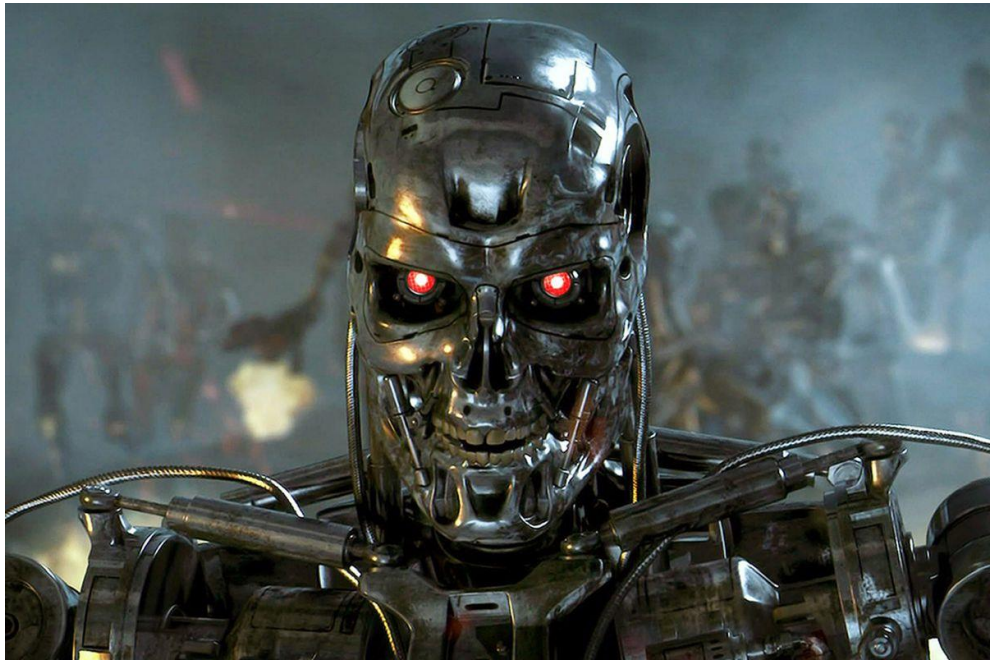

Incorrect Assumptions

- AI will automatically find patterns and violations with just data
  - **AI IS NOT MAGIC**
  - You have to give the data context
  - You have to have properly structured and collected data
  - You need to train your algorithms for several months on a large data set to get them right
  - You need well-documented use cases and scenarios to train your algorithms with for outcomes

Incorrect Assumptions

- AI is not this:

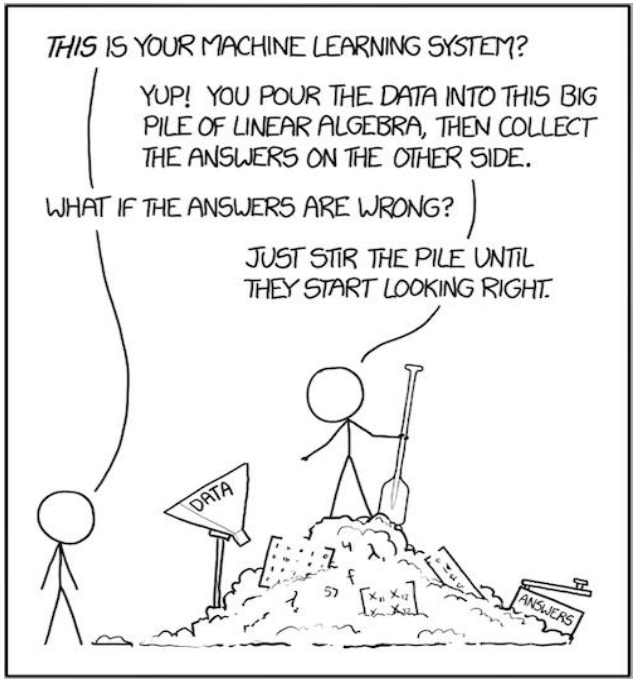

Risk Assessment

- Every year, we do a quantitative risk assessment
- We don’t have highs, mediums, or lows, we have scores
- Algorithm Score focuses on eight key areas:
  - Patient Safety
  - Reputational Risk
  - Potential Loss in Operating Income
  - Workforce Retention/Employee Engagement
  - Velocity of Risk
  - Likelihood
  - Patient Satisfaction
  - Potential Risk Contagion

Risk Assessment

- Our organization uses common risk scores
  - We all use the same scores and present top risks to the board
  - Our algorithm shows the magnitude of privacy and security issues across the organization
  - We made a conscious decision to use common language and sort high to low
    - We call it the 20/80 rule
    - Mitigating the top 20% of risks will also significantly mitigate the other 80% along the way if done right

Risk Assessment

- We don’t just use IS data
  - We also share significant information with Emergency Management
  - Our Business Impact Analysis program feeds into our IS Risk Management Program
    - If you don’t know what the business classifies as High Risk, then you’re not assessing properly
  - We are moving IU Health, IU School of Medicine, and IU Health Emergency Management to the same common platform for the 2019 assessments
  - Provide a better operating picture for the C-suite that shows actual risks

How the Journey Began

- It began with a need – to be able to meet the requirements for patient privacy monitoring
- Small Privacy team
- A very large task driven by the huge amount of events we accumulate
- Privacy Monitoring was driven on a reactive basis
- The need was demonstrated and quantified on our Risk Assessment

How the Journey Began

- Determine Desired Outcomes
  - Develop a significant breach deterrent
  - Demonstrate compliance with the HIPAA Security and Privacy Rules
  - An initial large amount of initial disciplinary cases that will trend downward
  - Process and Functionality Changes to address root causes of discovered issues
    - Expectation that this program would bring to light inefficiencies

Background on IU Health

- 17 Hospitals
- 34,100+ team members
- Primary teaching site for IU School of Medicine
- Largest health system in Indiana
- Combination of Academic Medical Centers, Community Hospitals, and Critical Access Hospitals
- Pediatric Level I Trauma Center at Riley Children's Health

Technical Background

- Interfacing with EMRs for audit logs is not easy
  - Even though both Epic and Cerner have made great strides (especially Cerner), it takes a lot of work to interface the systems
  - Many of the large vendors out there interface with Epic Clarity by default
- You also have to interface with your Human Capital Management System
  - You need to make sure that access is appropriate for team members
  - Under HITECH, you also need to account for unauthorized disclosures
  - You need to know demographics to be able to do this correctly

Technical Background

- Human Capital Management (HCM)
  - You also need to understand how your system handles job codes and job descriptions
  - This is key to identifying appropriate access
    - Back office (Rev Cycle, HIM, Finance) will have significant accesses
    - Quality and Regulatory will also have a lot
  - This work will also really help you on Payment Card Industry-Data Security Standards (PCI-DSS) compliance!

Starting with HR

- We met with our CHRO to discuss what we were doing and why
- We articulated the need to have this program in place and why
- We discussed that at no time AI would make a decision – a human analyst would
  - There is wariness of AI with top executives
  - They need to understand the algorithms and what they are looking for
  - There is a concern with false positives
- We discussed the cultural changes that would need to happen for this to be successful

HR Collaborative Meeting

- We then met with the HR Team
- As we are a large health system, the leadership team for HR is also correspondingly large
- We discovered that we needed to have unified policies for Corrective Action, along with our existing HIPAA, Acceptable Use, and Sanctions ones
- We addressed concerns about AI
- We discussed the desired outcomes

Senior Leadership Team

- With the assistance of several editors, including our Privacy Program Executive Director, Erica McDaniel, and CHIO, Dr. Seung Park, we presented to the senior leadership team
  - We discussed how the AI would work
    - No Computer Decisions – only results presented to a human analyst
    - Algorithms given data from only IU Health to train over several months before usage
    - Phased rollout of algorithms
    - Existing workflows and processes would be utilized
  - Extensive Communication Plan across enterprise in conjunction with Public Relations

Senior Leadership Team

- Dates given in advance
  - Including technical project plans
  - Including time to train AI
  - Significant testing time
  - All communication milestones
- Extensive list of teams given to SLT, CIO, and General Counsel that we would speak to

Organizational Staffing Plan

- How did we staff for this?
  - Vendor (Haystack) – Project Manager and Support Staff
  - Information Security Project Manager
  - Privacy Office – Executive Director
    - 2 Program Managers
    - 1 Staff Attorney
- We also had the Chief Privacy Officer and Information Security Officer available for any questions or issues
- Reports given to CIO and General Counsel

Policy Changes

- What policies do you need to change?
  - Corrective Action Policy – what happens when someone violates the rules
  - Sanctions Policy – required by the Security Rule for when someone violates the rules
  - Acceptable Disclosures Policy
  - Acceptable Computer Usage Policy
  - Organizational HIPAA Compliance Policy

Communication Plan for Leadership

- Who did we speak with outside the C-suite?
  - All physician leadership groups
  - All HR leadership groups
  - All Privacy/Security Liaisons across the organization
  - IU School of Medicine
  - Nursing Leadership
  - Revenue Cycle/Health Information Management
  - System Clinical Services
  - Chief Medical Officer Group
  - Hospital Medical Staffs
  - Affiliates
  - Information Systems Leaders

Communication Plan for Leadership

- We used the same template and presentation that we used for the C-suite – no special rules
- We made continual changes and corrections based on customer feedback
- We addressed key concerns with AI and ML
  - This was a recurring theme – customers wanted to know how the algorithms worked
- We got significant feedback from customers
  - Including ones who wanted to hook their EMR systems up to it
  - Most important, we communicated feedback and answered all questions!

Organizational Communication Plan

- We worked with PR to have planned communication using organizational channels
  - Intranet/Email/Department Newsletters
  - At least bi-weekly communications
- New Information Security Training Program
  - This came about due to Nursing expressing their dislike with the previous program
  - Actual incidents and scenarios in the training
  - Video based – we had minutes to complete it

Organizational Communication Plan

- We answered all questions asked – and we had many
- When we had go-live, we had a project manager and team monitoring the solution
- We had updates of progress to governance and our Privacy and Security Council
- Our major updates also had communication plans and meetings with the stakeholders
- We made sure to keep HR first and foremost with anything we did

Implementation and Monitoring

- We now have standard work built around daily monitoring of the system
  - While Cerner can give us data on an hourly basis, Epic has a daily feed from Clarity
  - We monitor the feeds from both systems
  - We have people who monitor the system every morning for new potential issues and cases
  - We have integrated this into our workflow

Continual Follow up and Metrics

- We have monthly operational reporting on potential cases, investigations, actual cases, and reported breaches
- Reporting goes to both CIO and General Counsel
- Monthly reporting to Governance committees
- Our board also gets reports on cases

How has this influenced future work for us?

- Haystack was the first major non-clinical and non-research AI/ML project we put into place that had the attention of senior leadership
- We learned several important lessons:
  - Constant, clear communication about what we are doing with AI/ML
  - Explain the algorithms and data used
  - Explain the workflows and processes in detail
  - Make sure that our customers understand that the AI does not make a final decision for us

What did we do afterward?

- We implemented contact language and guidance on AI/ML/Deep Learning/Robotic Process Automation
- We worked with Regenstrief Institute on this
- Our language focuses on:
  - Use Cases
  - Data Provenance and Appropriate Consent
  - Minimum Necessary Data
  - Verification and Validation of Algorithms
  - Explanation of Algorithms
  - Potential use of Blockchain/DLT to validate and verify data

What do we expect?

- This first use case provided us an excellent plan to implement future solutions
- We were able to demonstrate that we understand AI and how to properly utilize it
- We expect to bring on multiple future solutions across the organization using the lessons we learned
  - Not reinvent the wheel each time
  - Expect an outgrowth of AI apps
  - Want to avoid the hype of AI through experience

Thank you!

- Questions?
- Follow me on Twitter at @mitchparkerciso

A

A