IEEE/ACM CHASE 2024

Systematically Assessing the Security Risks of AI/ML-enabled Connected Healthcare Systems


Mohammed Elnawawy*


Mohammadreza Hallajiyan*


Gargi Mitra*


Shahrear Iqbal$


Karthik Pattabiraman*

*University of British Columbia (UBC), $National Research Council, Canada

Why focus on security of ML-enabled medical devices?

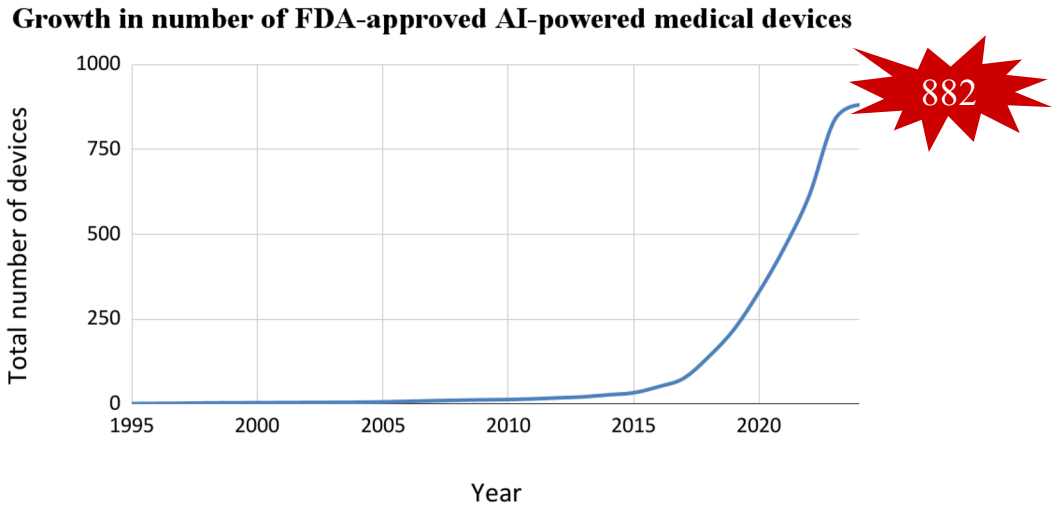

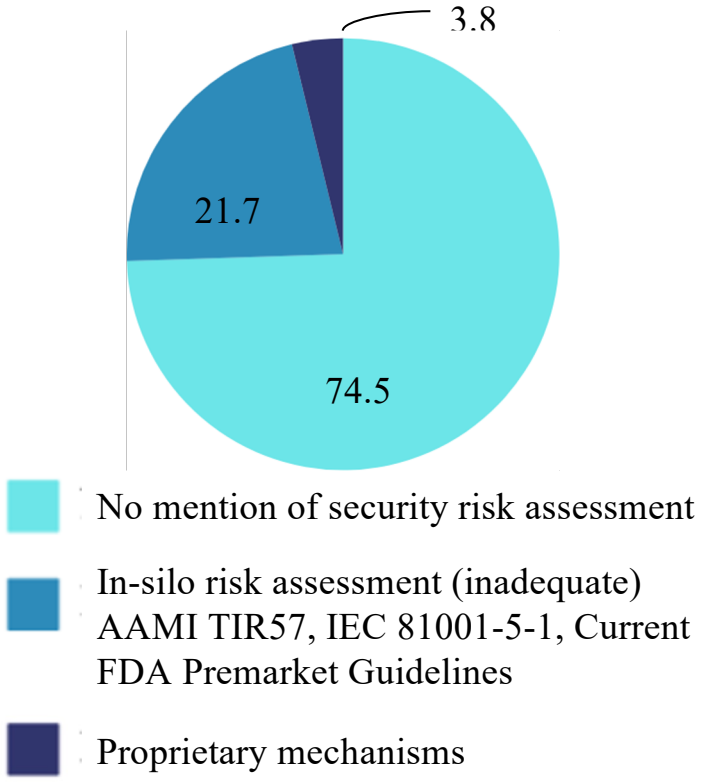

Recent FDA guidelines

- Pre-market security assessment: Mandatory
- Design for Security with no implicit trust

Source: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices, Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions by FDA

Why is securing **ML-enabled** medical devices challenging?

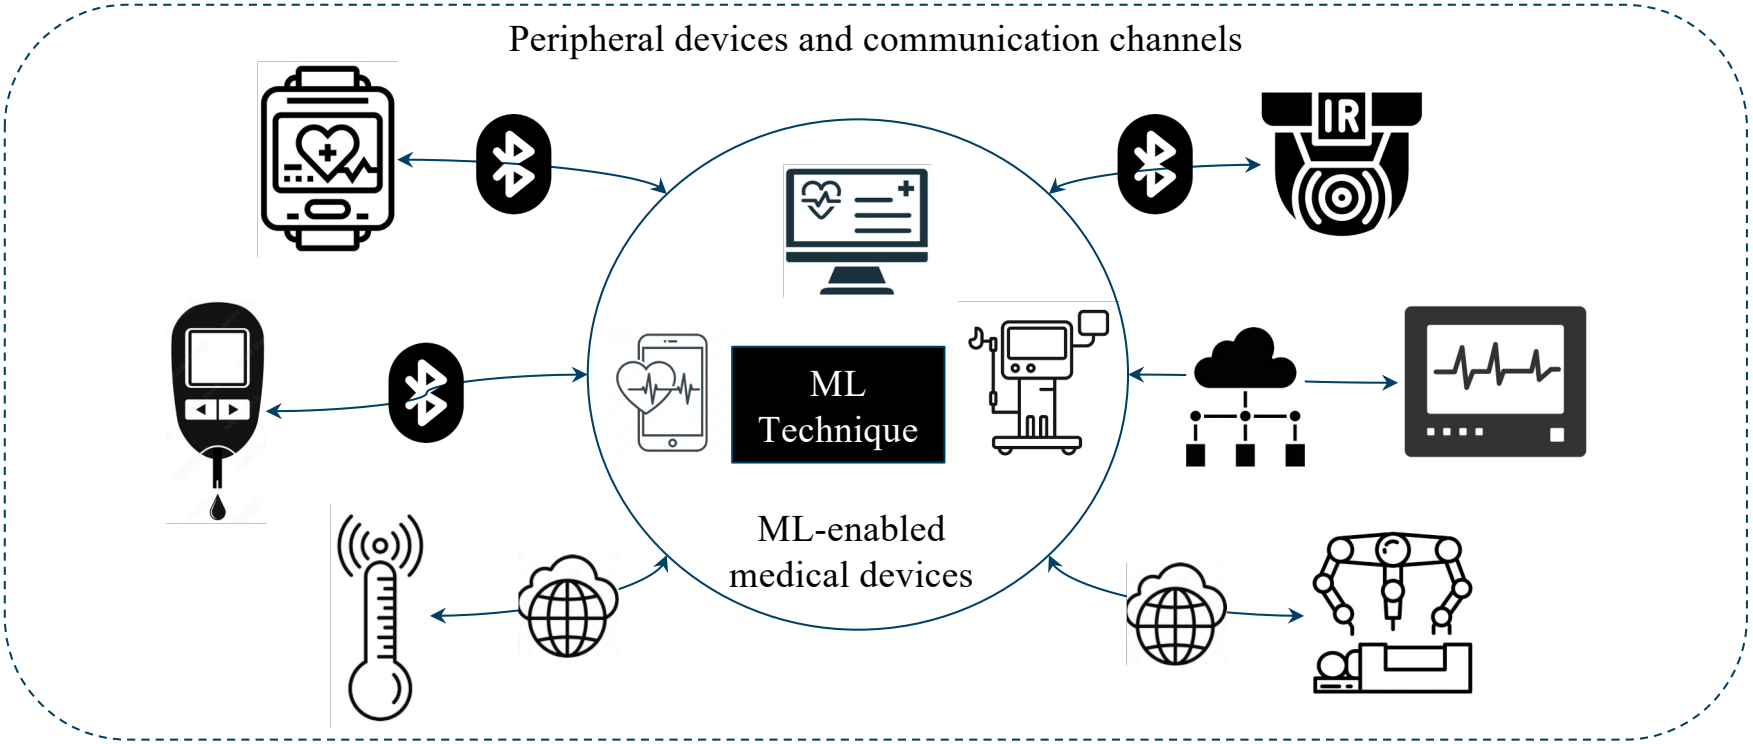

Highly interconnected multi-vendor setup

Why is securing **ML-enabled** medical devices challenging?

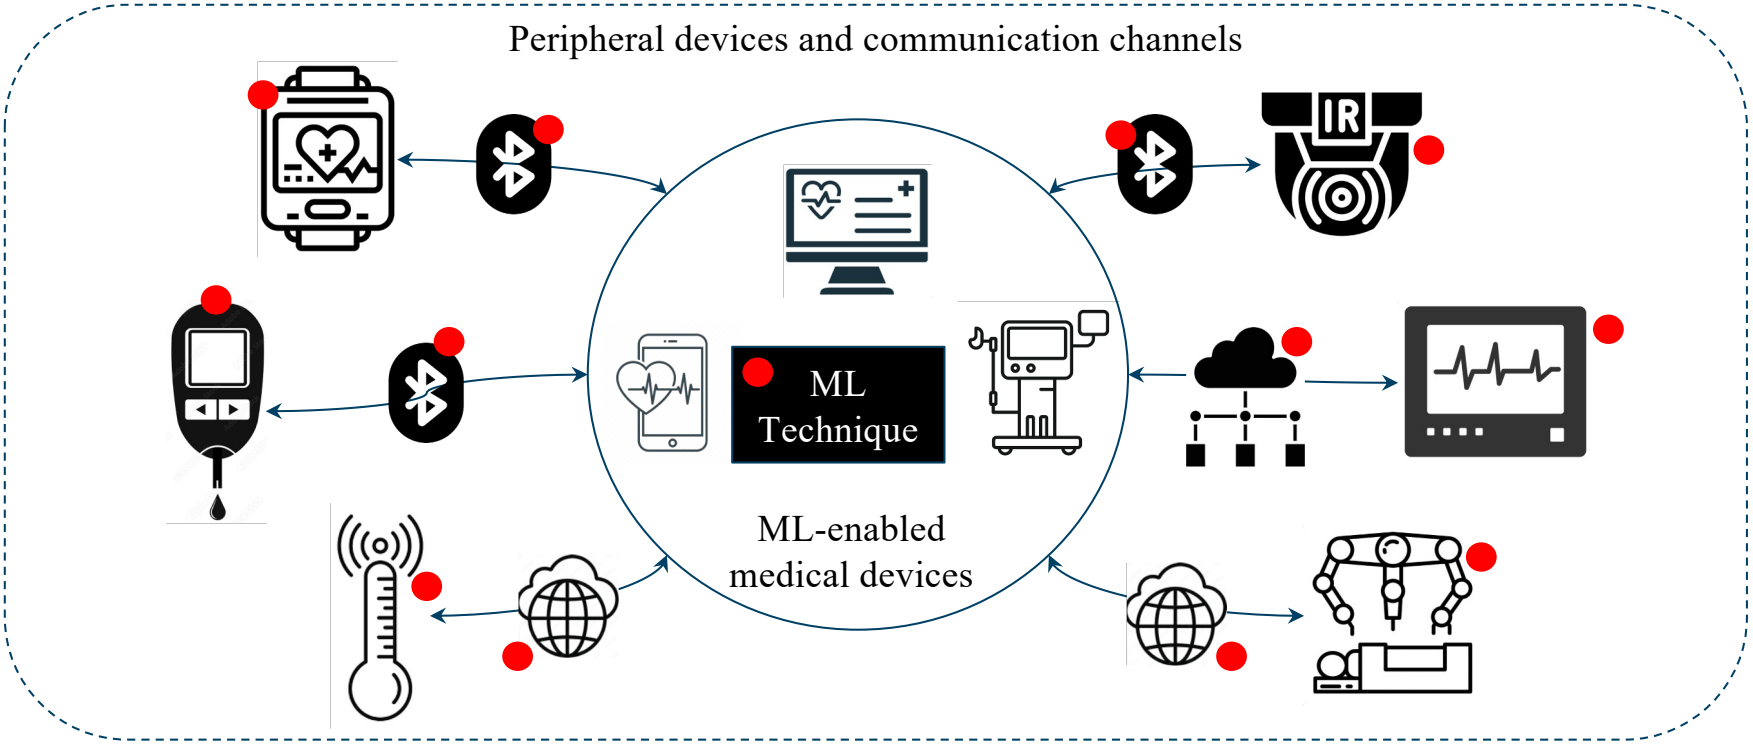

Large number of attack points (attack surface) - Hard to foresee during design

What can go wrong ?

A short story inspired by experiments

ML-enabled Blood Glucose Management System (BGMS)

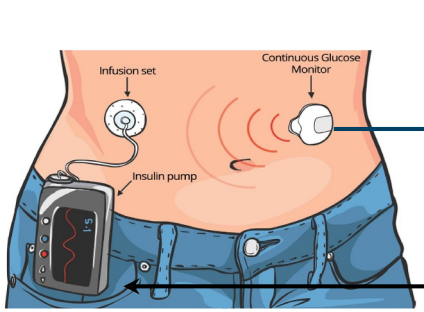

Patient

Blood glucose reading

Insulin dose

ML-based Diabetes management app

Meal timing, Carbs taken, Age, Gender

Patient

ML-enabled Blood Glucose Management System (BGMS)

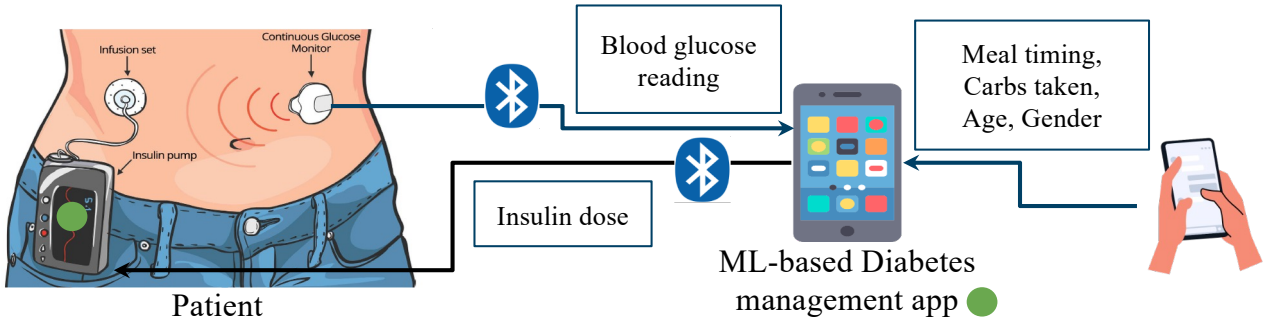

Here's the latest AI-powered insulin pump!

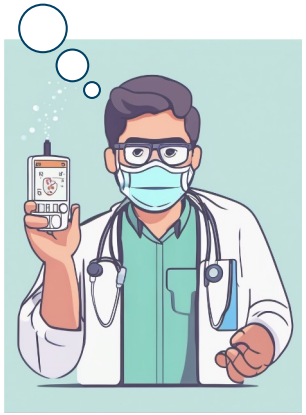

99.9% accurate

safe ✓

secure ✓

Compatible with many models of glucose meters

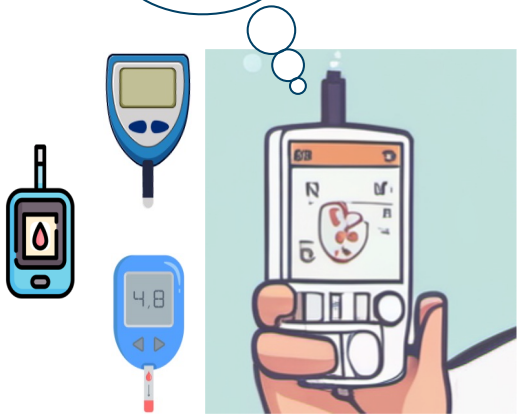

I want the new AI-powered insulin pump!

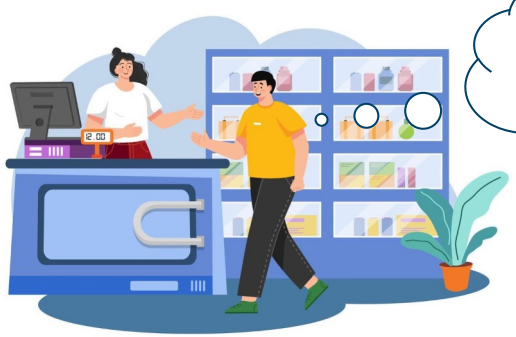

I'll get **Glucose meter A**

I'll buy **Glucose meter B**

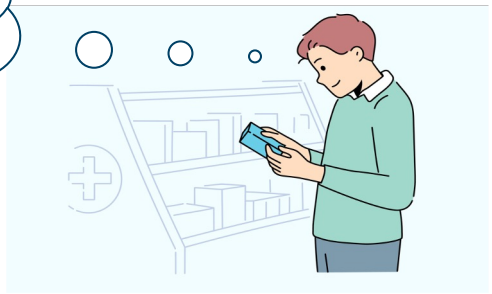

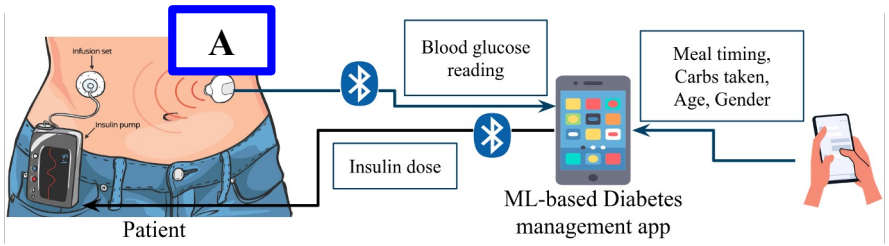

Patient #1

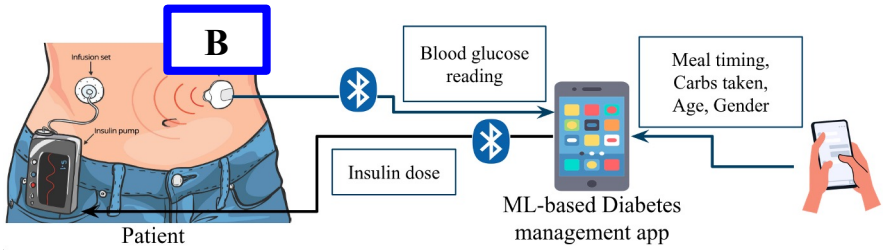

Patient #2

Let's try evasion attack!

Let's make the app suggest a HIGH insulin dose when he DOESN'T need it

But how do I inject fake readings into the system?

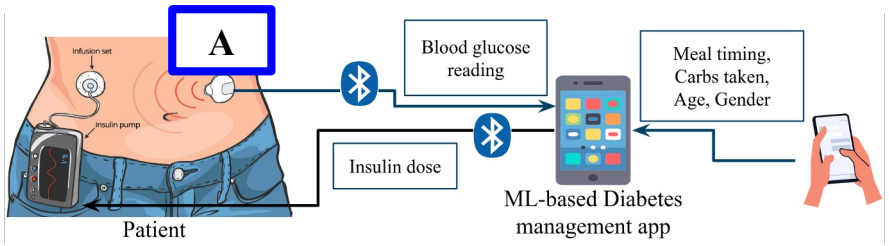

Patient #1

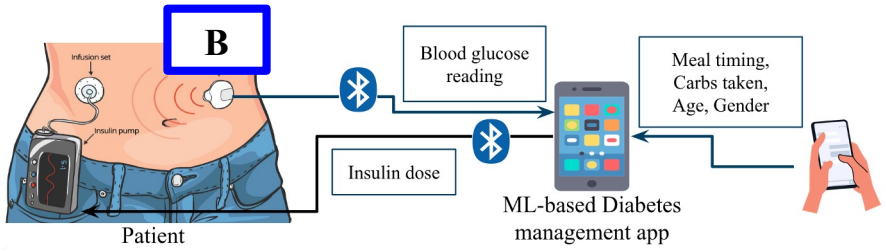

Patient #2

Let's try evasion attack!

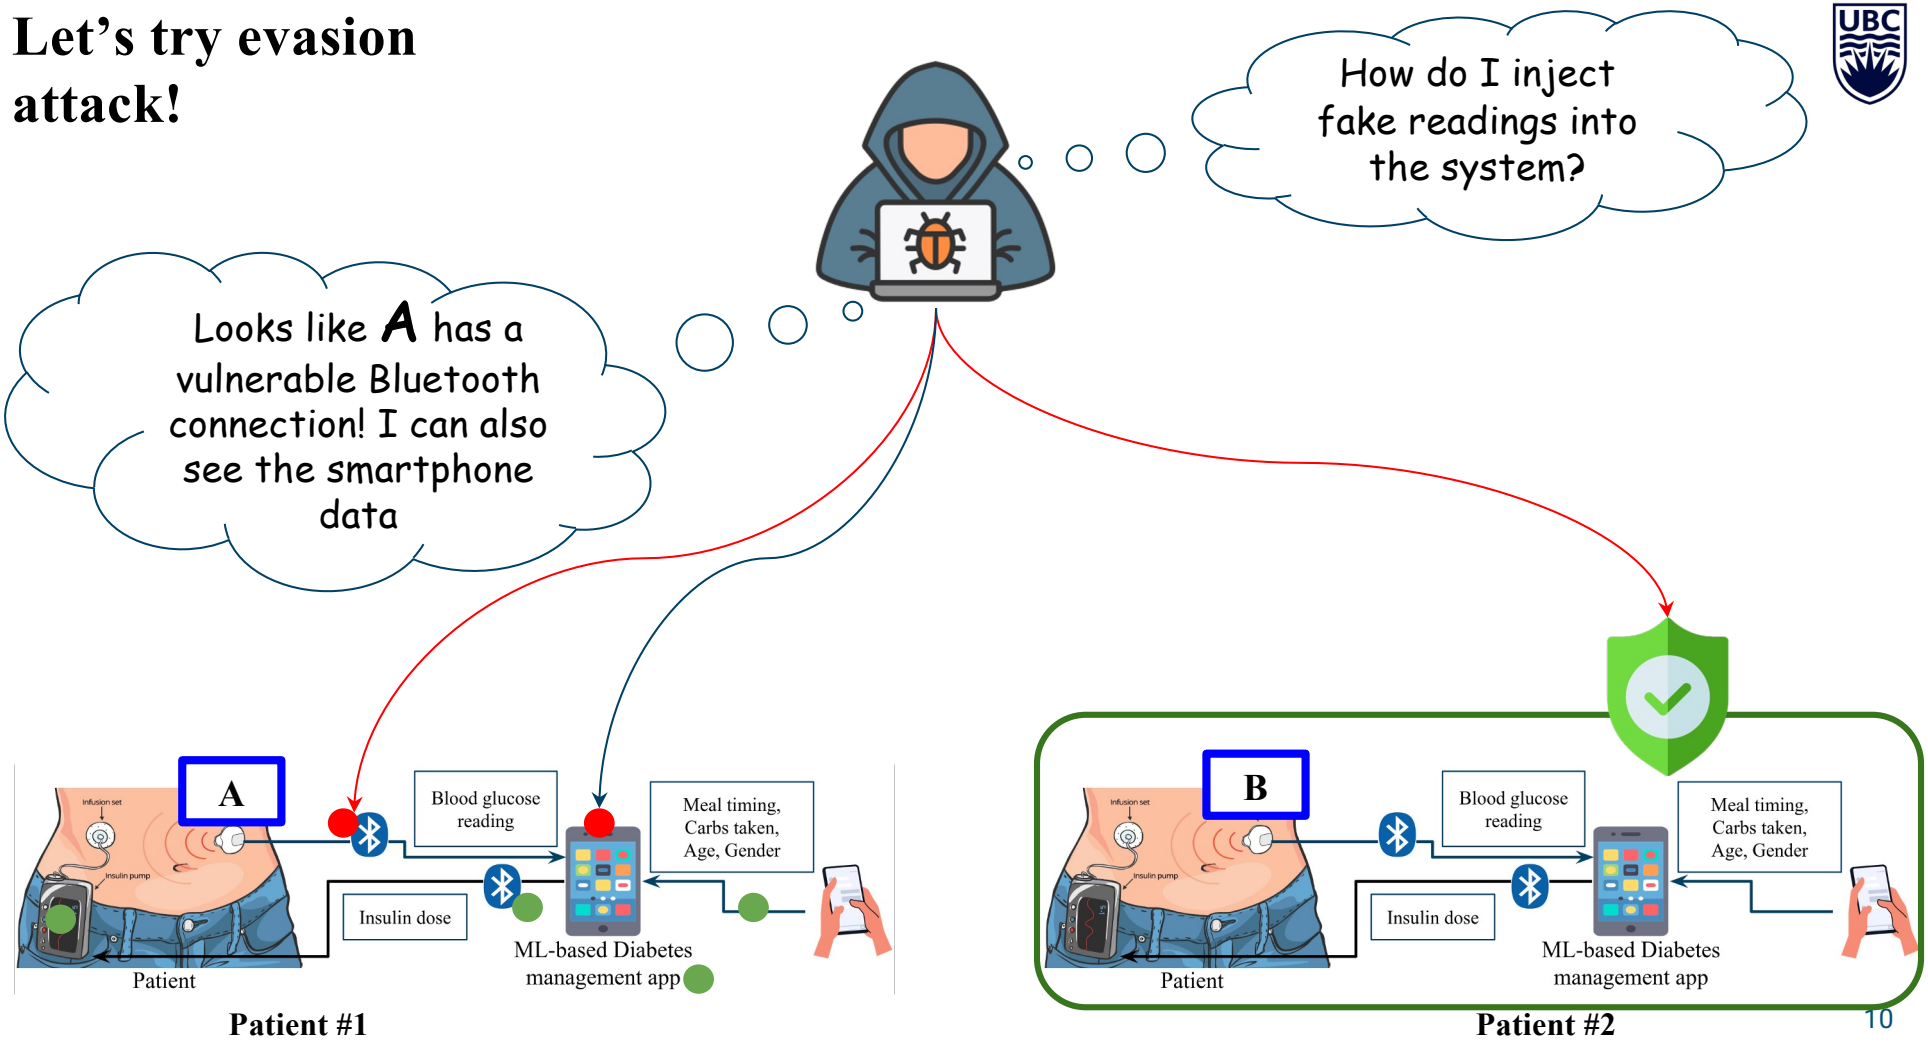

Let's try evasion attack!

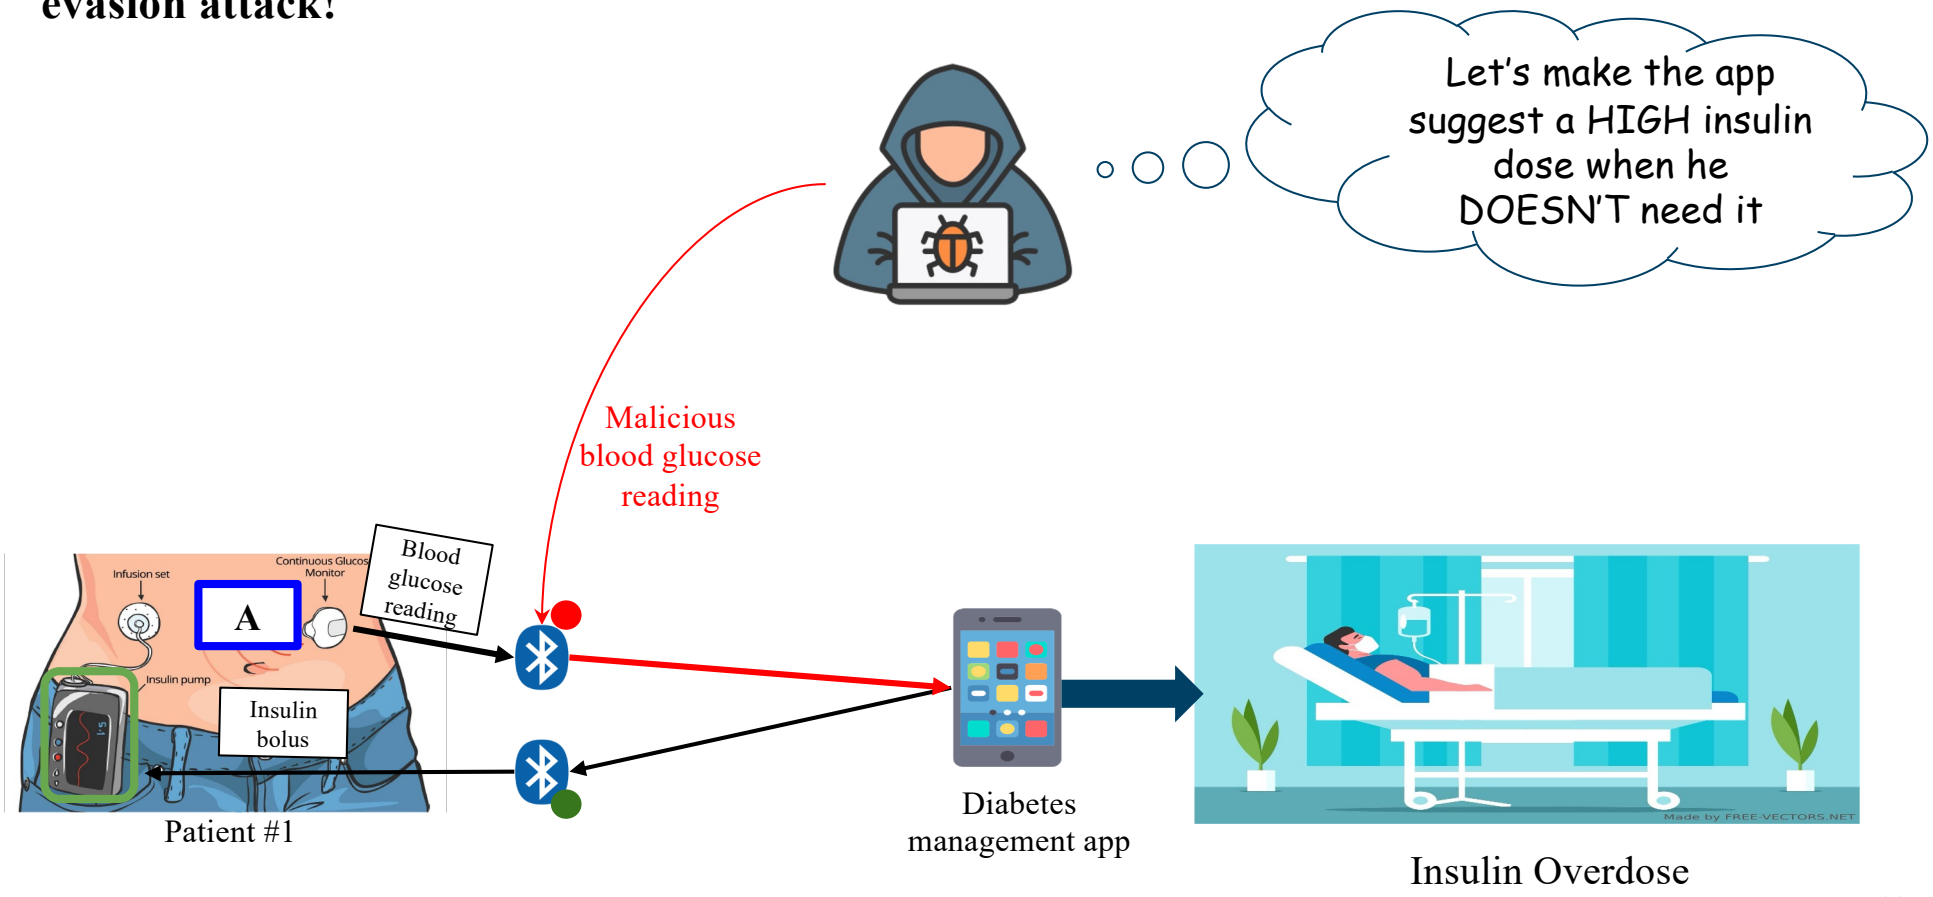

Our Case Study

Let's try evasion attack!

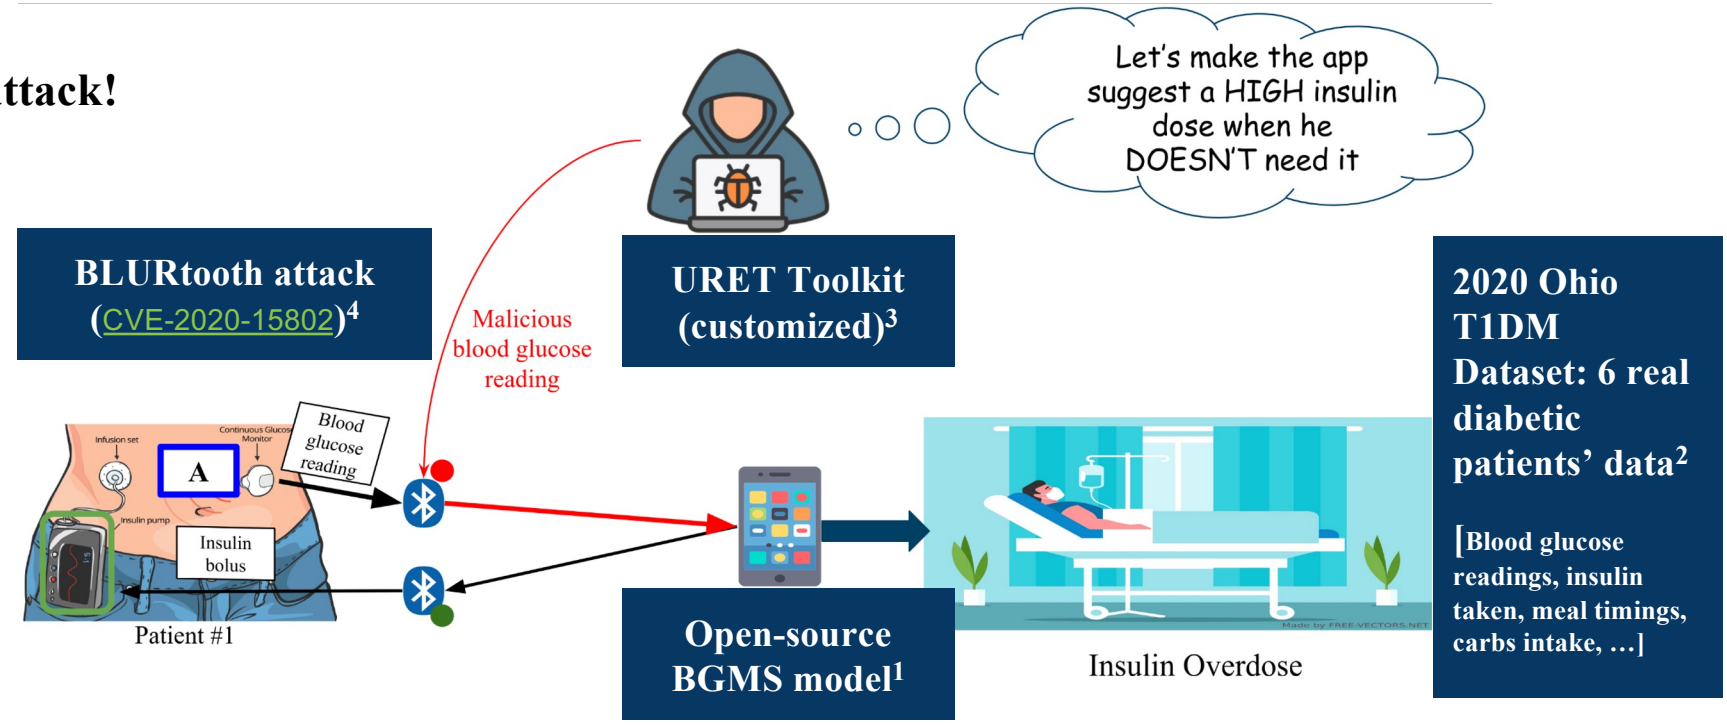

[1] Harry Rubin-Falcone, Ian Fox, and Jenna Wiens. **Deep Residual Time-Series Forecasting: Application to Blood Glucose Prediction**. In KDH@ECAI, 2020.
[2] Cindy Marling and Razvan Bunescu. **The OhioT1DM dataset for blood glucose level prediction**. In CEUR workshop proceedings, NIH Public Access, 2020.
[3] Kevin Eykholt, Taesung Lee, Douglas Schales, Jiyong Jang, and Ian Molloy. **URET: Universal Robustness Evaluation Toolkit (for Evasion)**. In USENIX Security 2023
[4] Kasper Rasmussen. **BLURtooth: Exploiting Cross-Transport Key Derivation in Bluetooth Classic and Bluetooth Low Energy**. In AsiaCCS, 2022.

How can manufacturers foresee post-deployment security risks?

A systematic assessment of 20 FDA-approved devices

Security Assessment Process

Security Assessment Process

1. ML technique used by the device

- Device functionality
- ML engine mispredicts: Impact on patient?
- Input data types
- ML technique

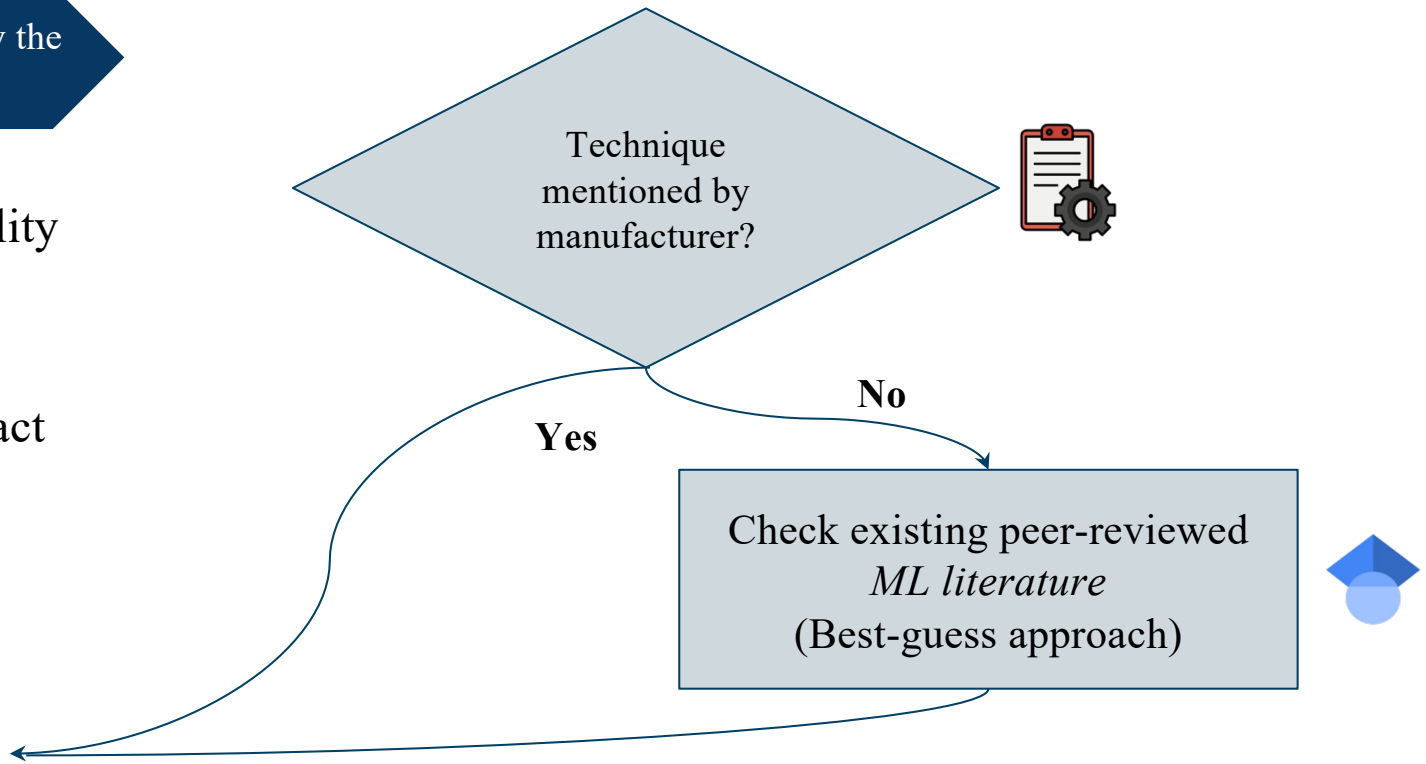

Security Assessment Process

1. ML technique used by device
2. Vulnerability to inference-time attacks

ML technique in (1)

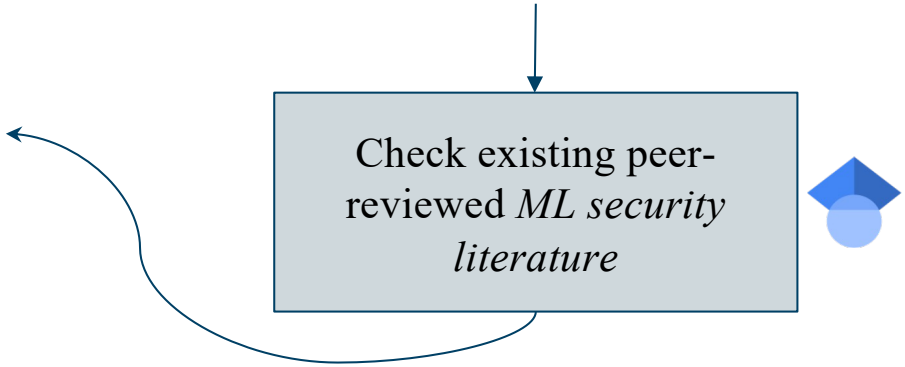

- Known attacks on ML technique
- Checks performed to detect malicious inputs? **No**, unless mentioned by manufacturer

Security Assessment Process

Device description in (1)

- Compatible peripheral devices and communication protocols

Security Assessment Process

1. ML technique used by the device
2. Vulnerability to inference-time attacks
3. Peripheral device compatibility
4. Peripheral device vulnerability

Attack identified in (2),
Peripheral devices identified in (3)

Does any device / combination of devices have vulnerability that allows an attacker to execute the attack?

- Vulnerable devices and vulnerability descriptions
- Attacker position and capabilities

Devices Assessed

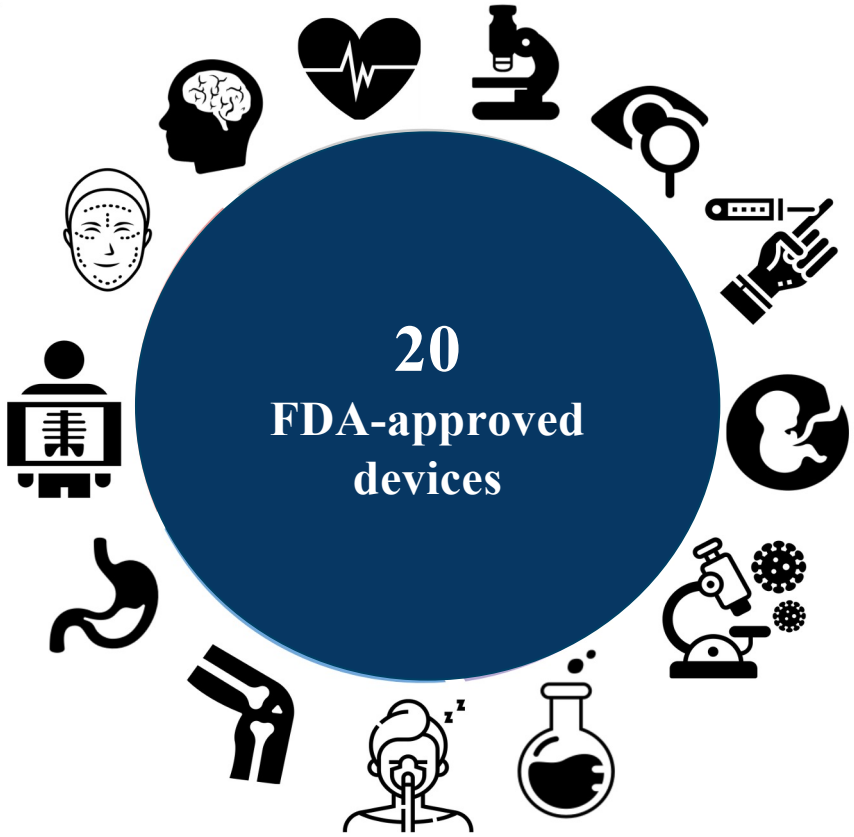

SaMD

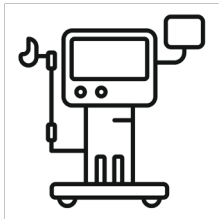

SiMD

Self-monitoring
patients

Used in
hospitals

Interesting insights from the assessment

1. **Post-deployment attacks :** 16/20 devices vulnerable
2. **Attack Surface:** SaMD > SiMD
3. **Widespread Vulnerabilities :** Attack point - Core technology (e.g., IR cameras)
4. **Hard-to-detect attack paths:** e.g., IDx-DR software.
5. **Persistent Vulnerabilities:** Some won't be fixed by OEMs.

Summary

- **ML-enabled medical applications:** Large, complex attack surface → Health risk
- **Our contribution:**
  - Systematic end-to-end security assessment process
  - Case study demonstration
- **Next steps:**
  - Automate assessment technique
  - Profile patients by security risk

Systematically Assessing the Security Risks of AI/ML-enabled Connected Healthcare Systems, *Mohammad ElNawawy, Mohammadreza Hallajiyan, Gargi Mitra, Shahrear Iqbal, and Karthik Pattabiraman,* IEEE/ACM international conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE), 2024