Artificial Intelligence + *Industrial Control Systems*

Adopting AI to Protect ICS: Assessing Challenges and Opportunities

From the operators' perspective!

Clement Fung, Eric Zeng, Lujo Bauer
Carnegie Mellon University

S3D

CyLab Carnegie Mellon University Security and Privacy Institute

What are industrial control systems (ICS)?

- ICS are systems that control processes in critical infrastructure

[1] https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors

What are industrial control systems (ICS)?

- ICS are systems that control processes in critical infrastructure

Energy

Water Treatment

Manufacturing

[1] https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors

ICS are common targets for attack

ICS are common targets for attack

BlackEnergy (2015)
Industroyer (2016)

ICS are common targets for attack

BlackEnergy (2015)
Industroyer (2016)

Colonial Pipeline (2021)

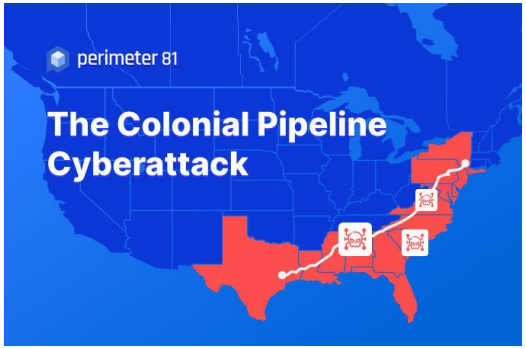

ICS are common targets for attack

BlackEnergy (2015)
Industroyer (2016)

Colonial Pipeline (2021)

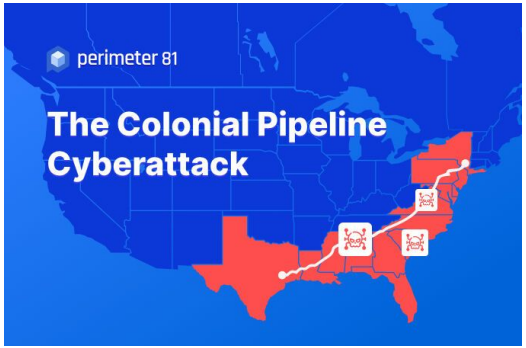

Aliquippa Water Plant (2023)

ICS are common targets for attack

BlackEnergy (2015)
Industroyer (2016)

Colonial Pipeline (2021)

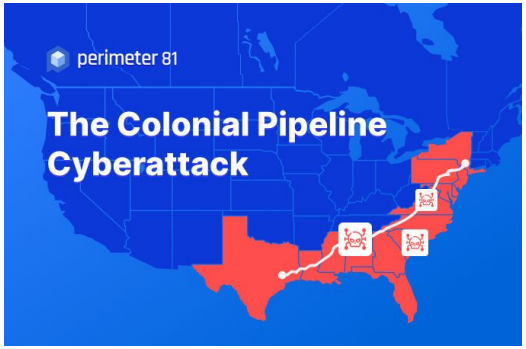

Aliquippa Water Plant (2023)

- July 2021: US President issues National Security Memorandum
  - “Improving Cybersecurity for Critical Infrastructure Control Systems”

Protecting ICS: current practice

Industrial process

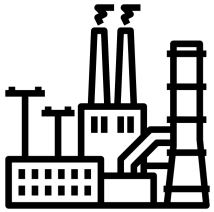

Protecting ICS: current practice

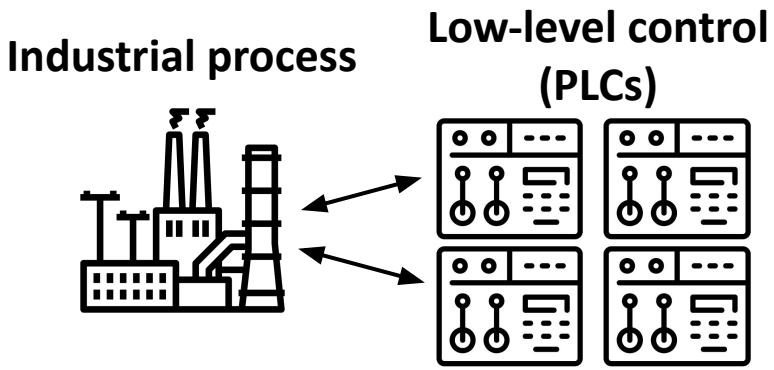

Protecting ICS: current practice

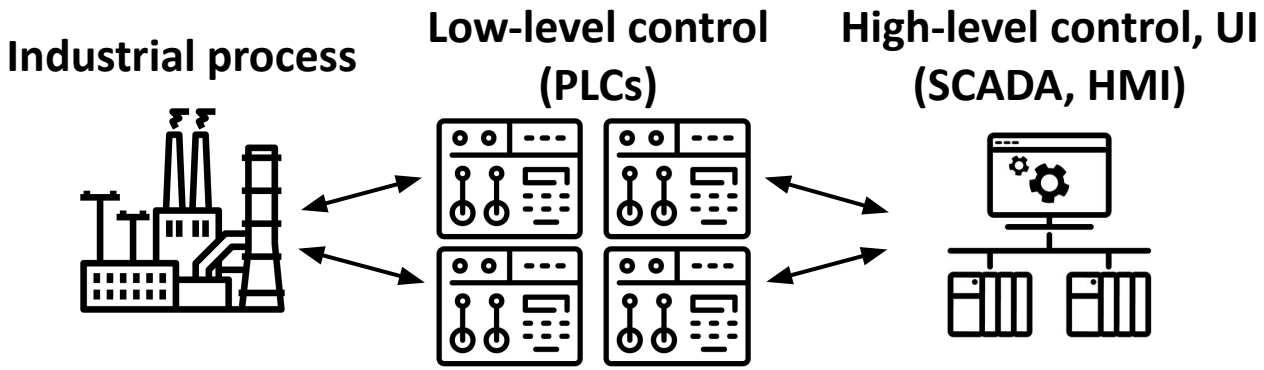

Protecting ICS: current practice

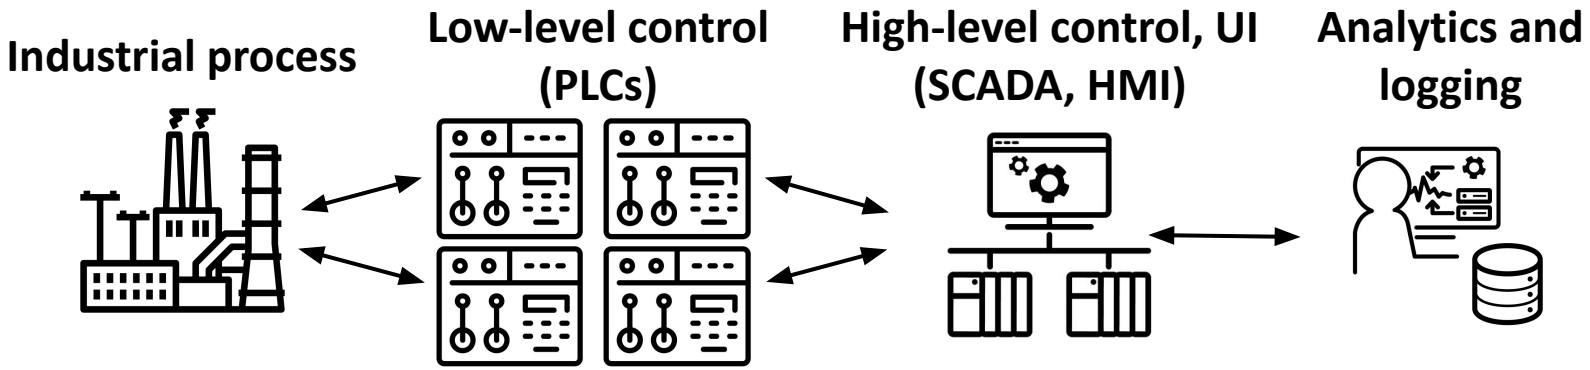

Protecting ICS: current practice

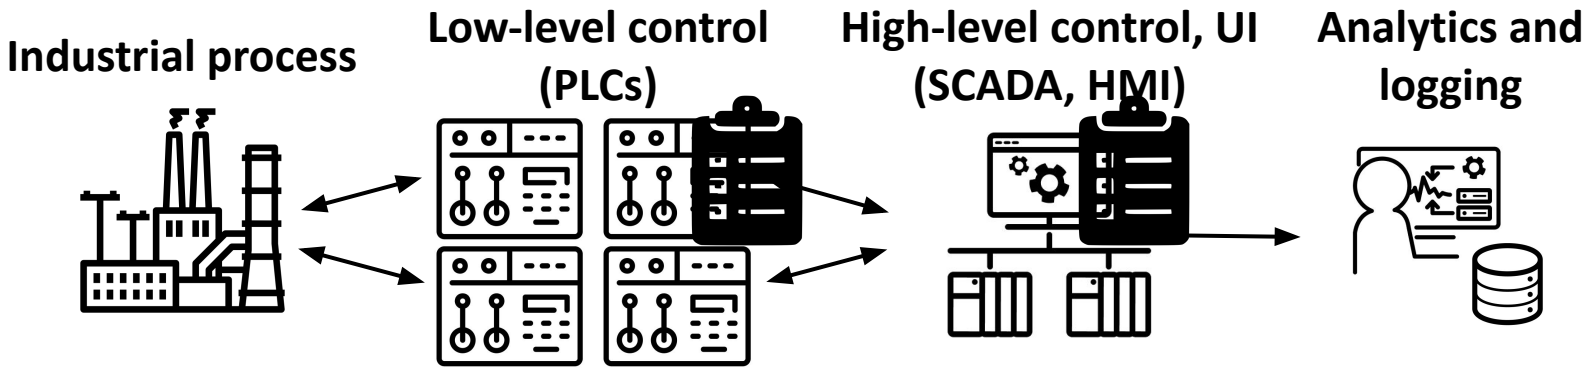

ICS security in practice:

- Manually write detection rules

Protecting ICS: current practice

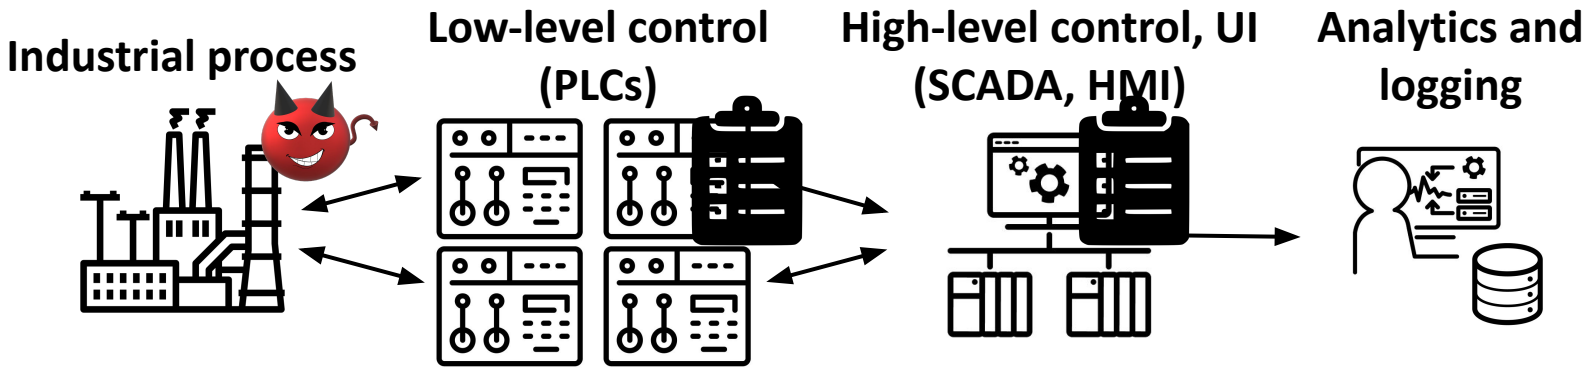

ICS security in practice:

- Manually write detection rules

Protecting ICS: current practice

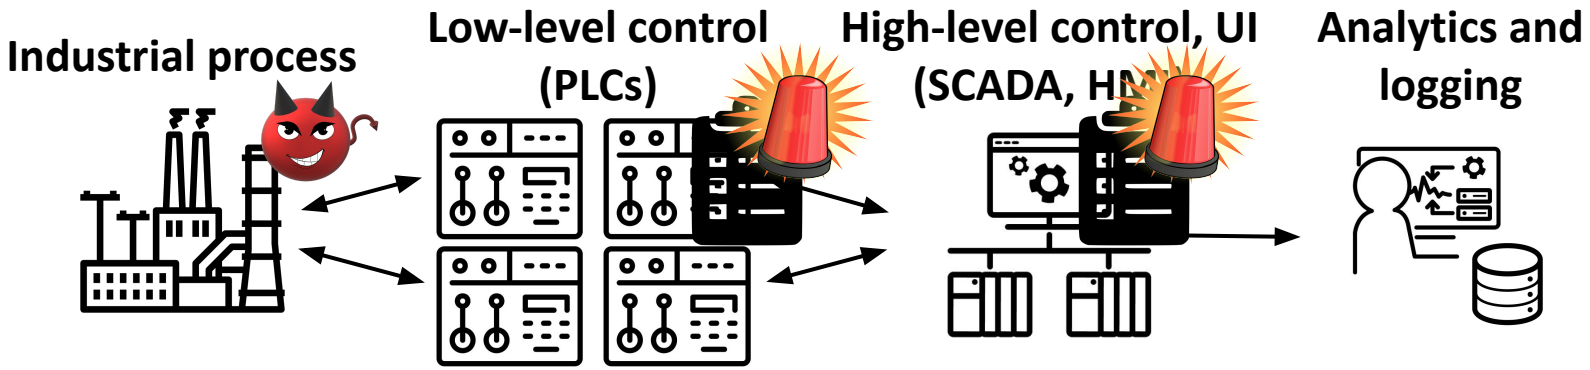

ICS security in practice:

- Manually write detection rules

Protecting ICS: current practice

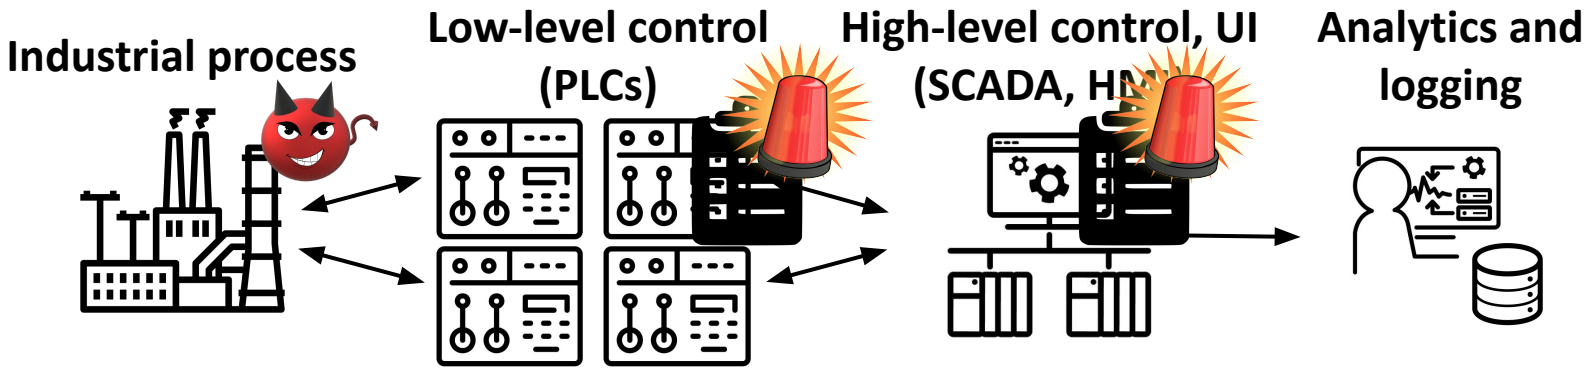

ICS security in practice:

- Manually write detection rules
- **Imperfect** and **labor-intensive**

Protecting ICS: in research, using AI

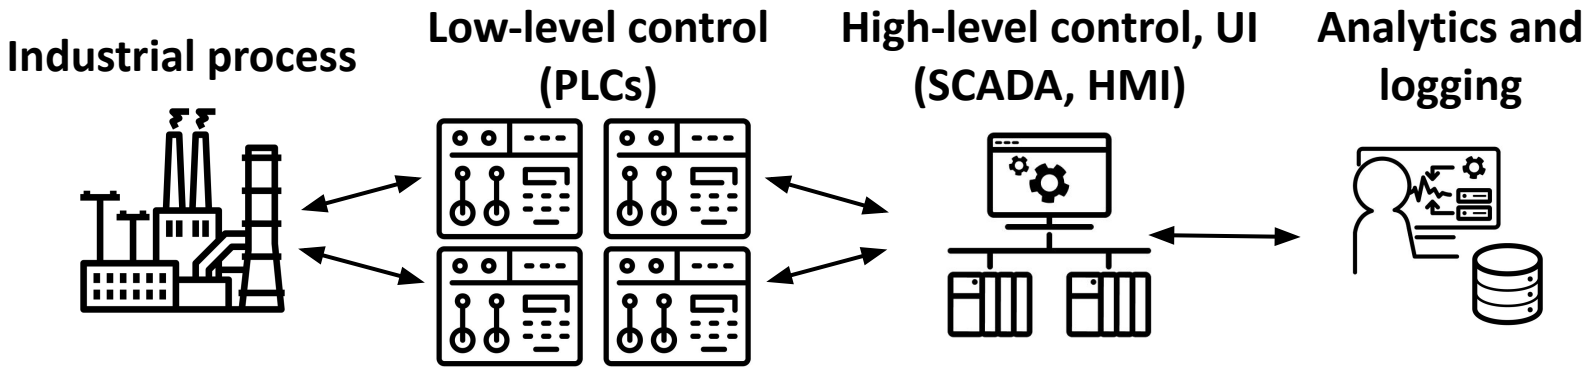

Protecting ICS: in research, using AI

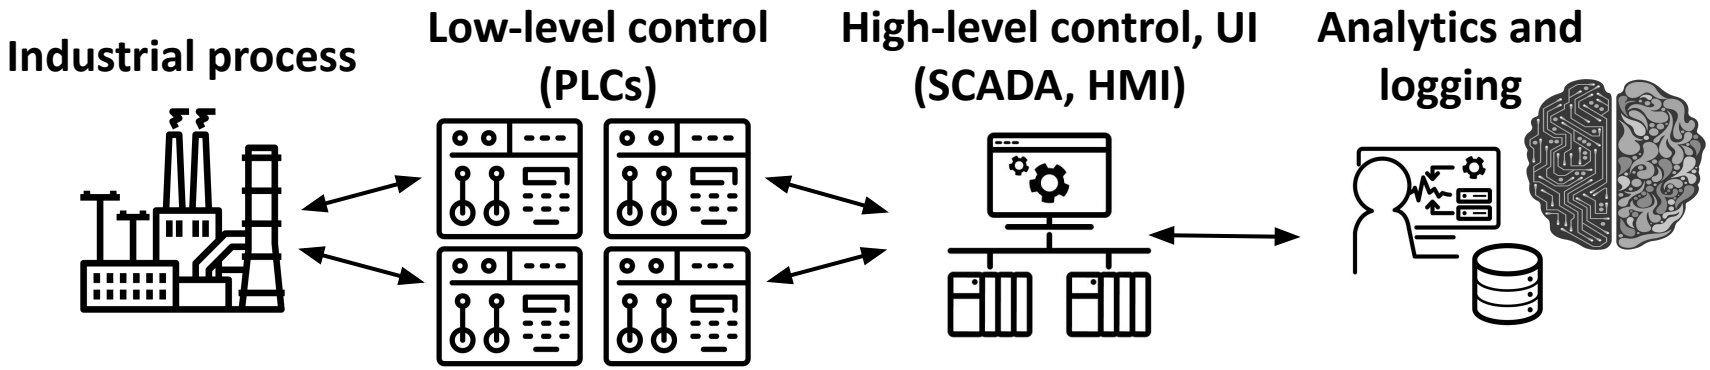

AI for ICS security in **research papers**:

- Train an AI model on ICS data

Protecting ICS: in research, using AI

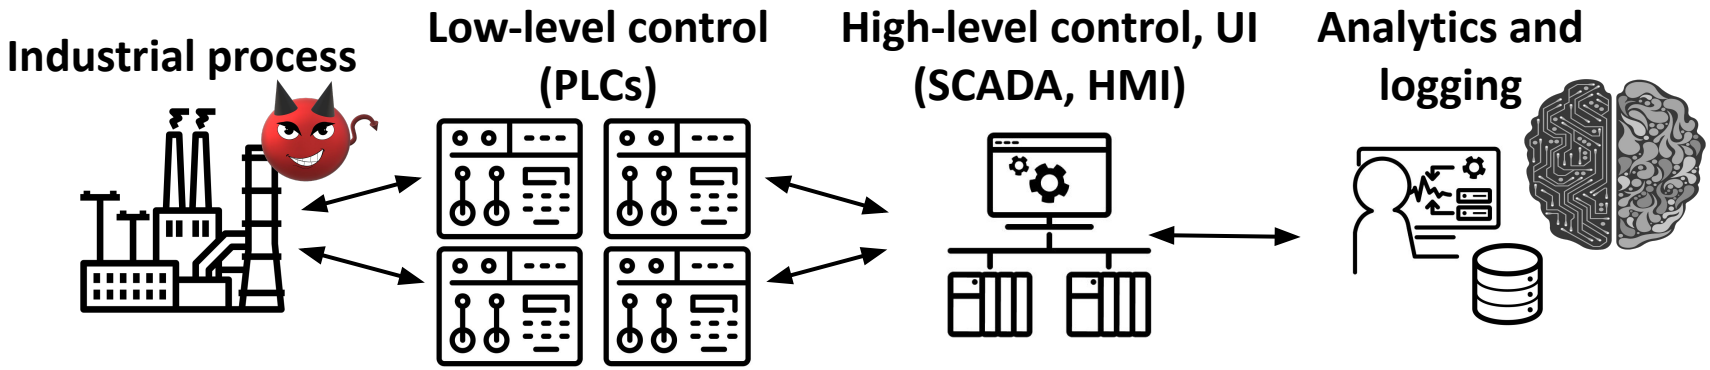

AI for ICS security in **research papers**:

- Train an AI model on ICS data

Protecting ICS: in research, using AI

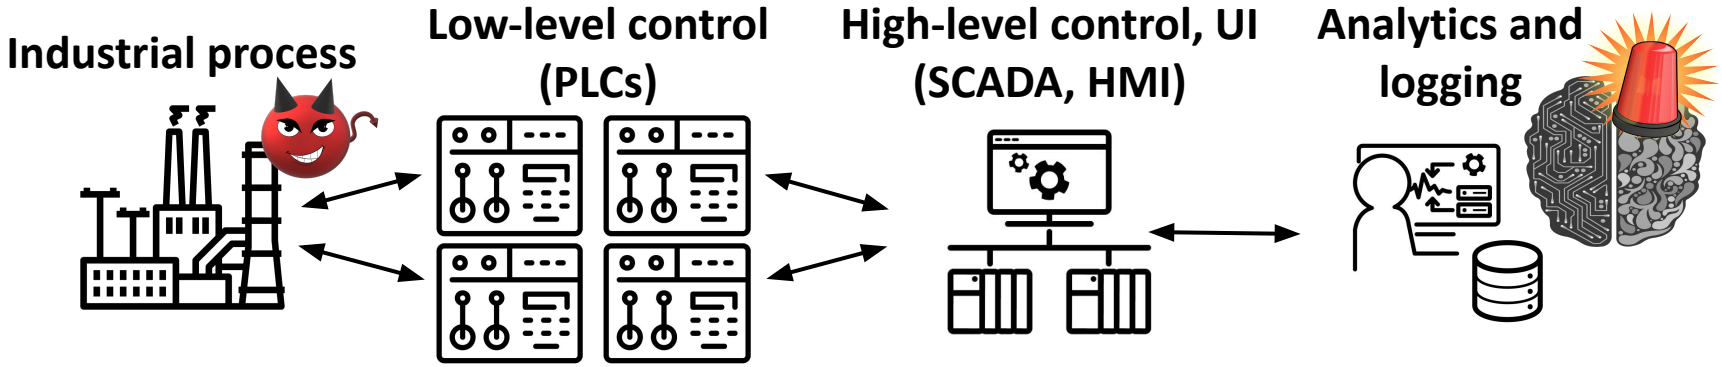

AI for ICS security in **research papers**:

- Train an AI model on ICS data
- Detect attacks with high accuracy

Protecting ICS: in research, using AI

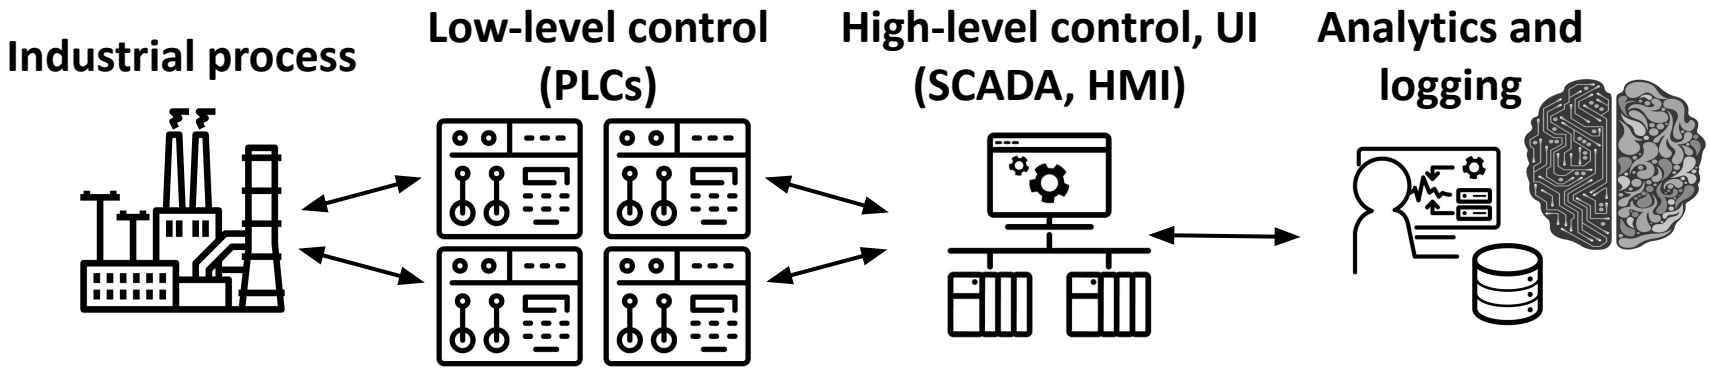

AI for ICS security **in practice** (2024):

Protecting ICS: in research, using AI

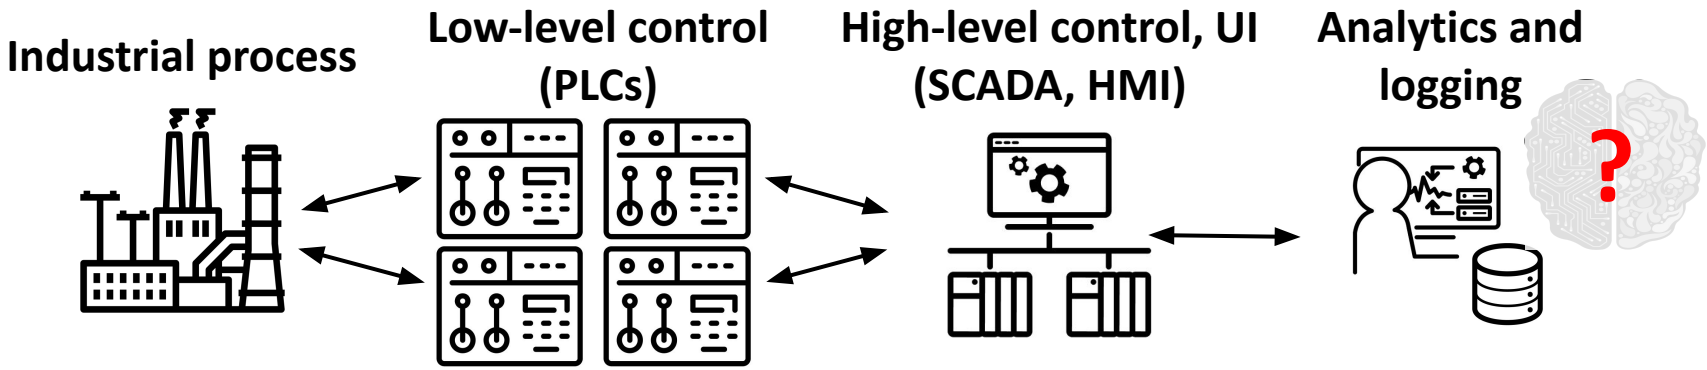

AI for ICS security **in practice** (2024):

- 10% are using AI in ICS networks
- 19% are experimenting with AI

Protecting ICS: in research, using AI

Industrial process

Low-level control (PLCs)

High-level control, UI (SCADA, HMI)

Analytics and logging

In this work, we investigate this gap between research and practice:

- 10% are using AI in ICS networks
- 19% are experimenting with AI

Protecting ICS: in research, using AI

Industrial process

Low-level control (PLCs)

High-level control, UI (SCADA, HMI)

Analytics and logging

In this work, we investigate this gap between research and practice:

We **interview practitioners who work on protecting ICS** to understand their practices, pain points, and requirements

- 10% are using AI in ICS networks
- 19% are experimenting with AI

Who is involved with protecting ICS?

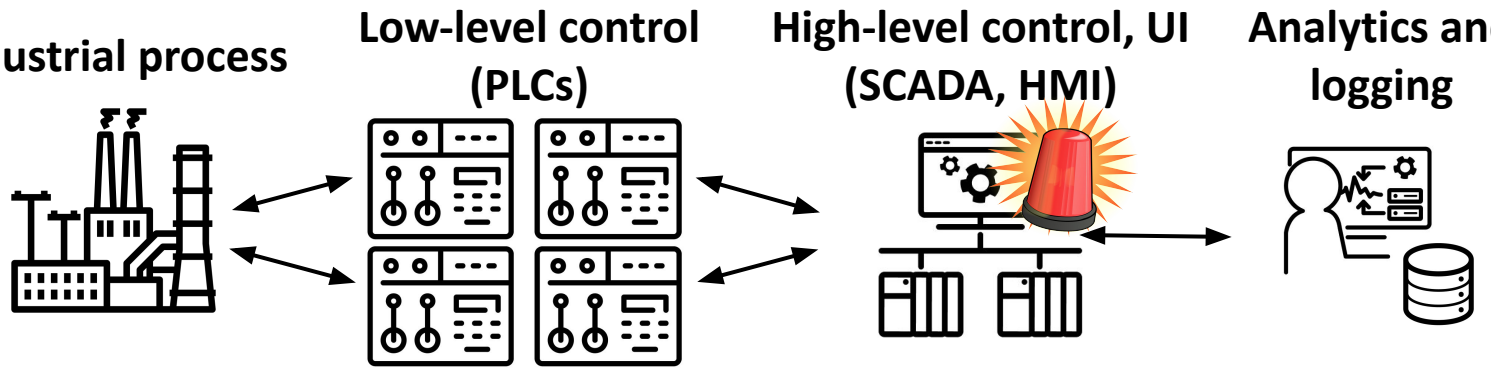

Technology

Humans

Who is involved with protecting ICS?

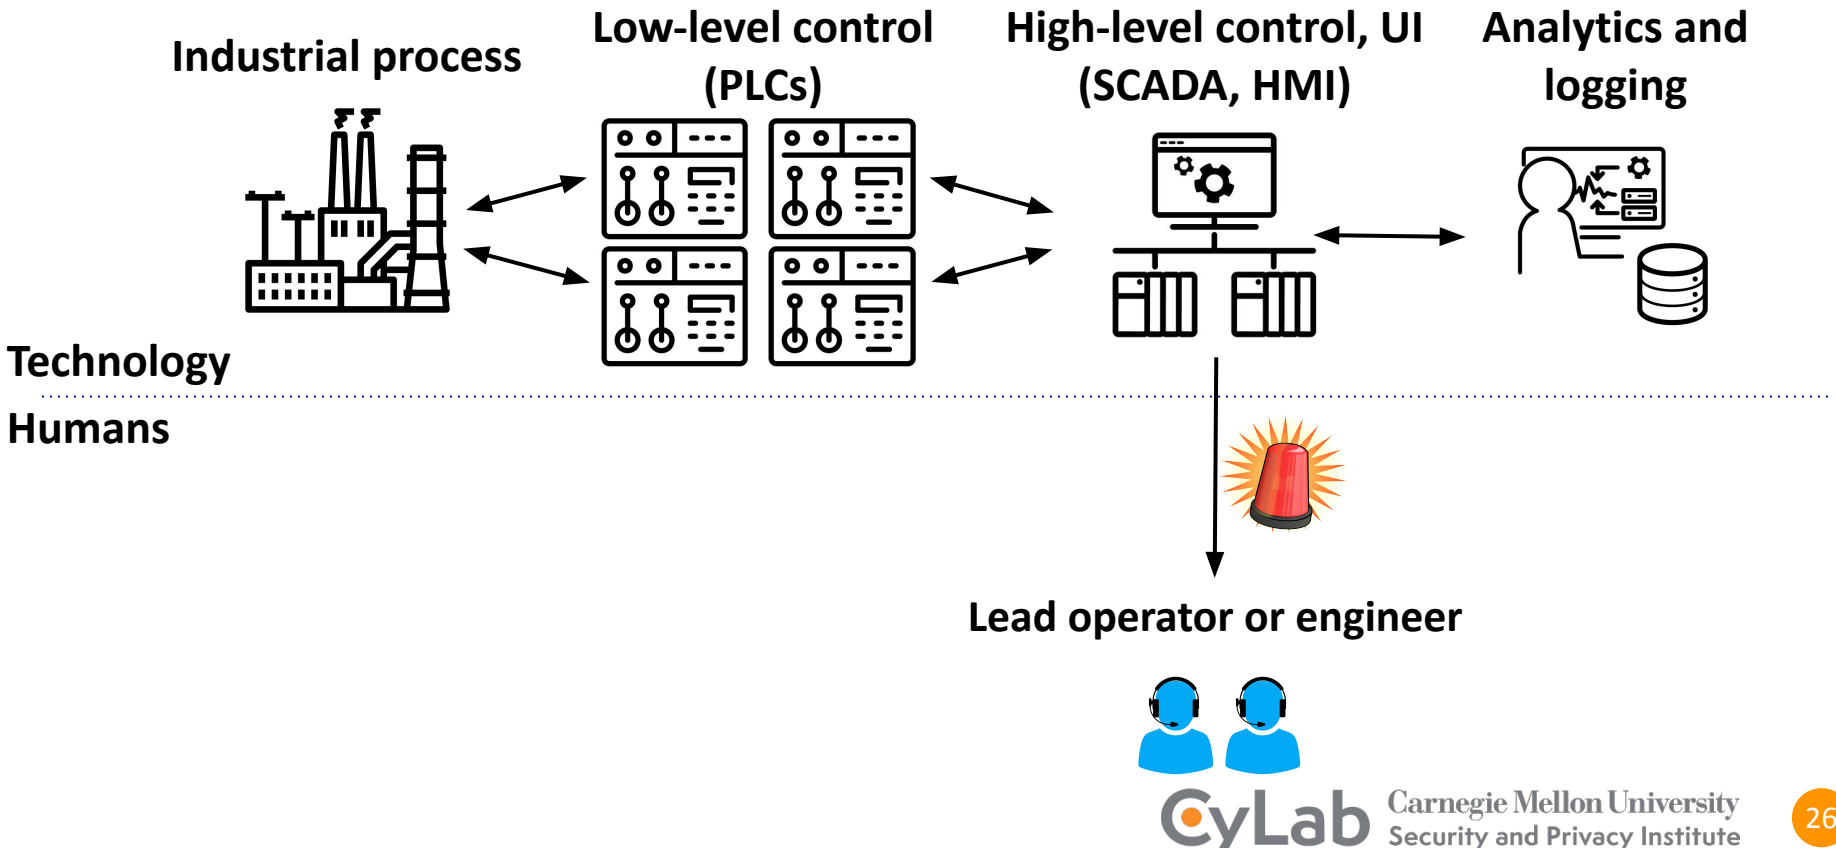

Who is involved with protecting ICS?

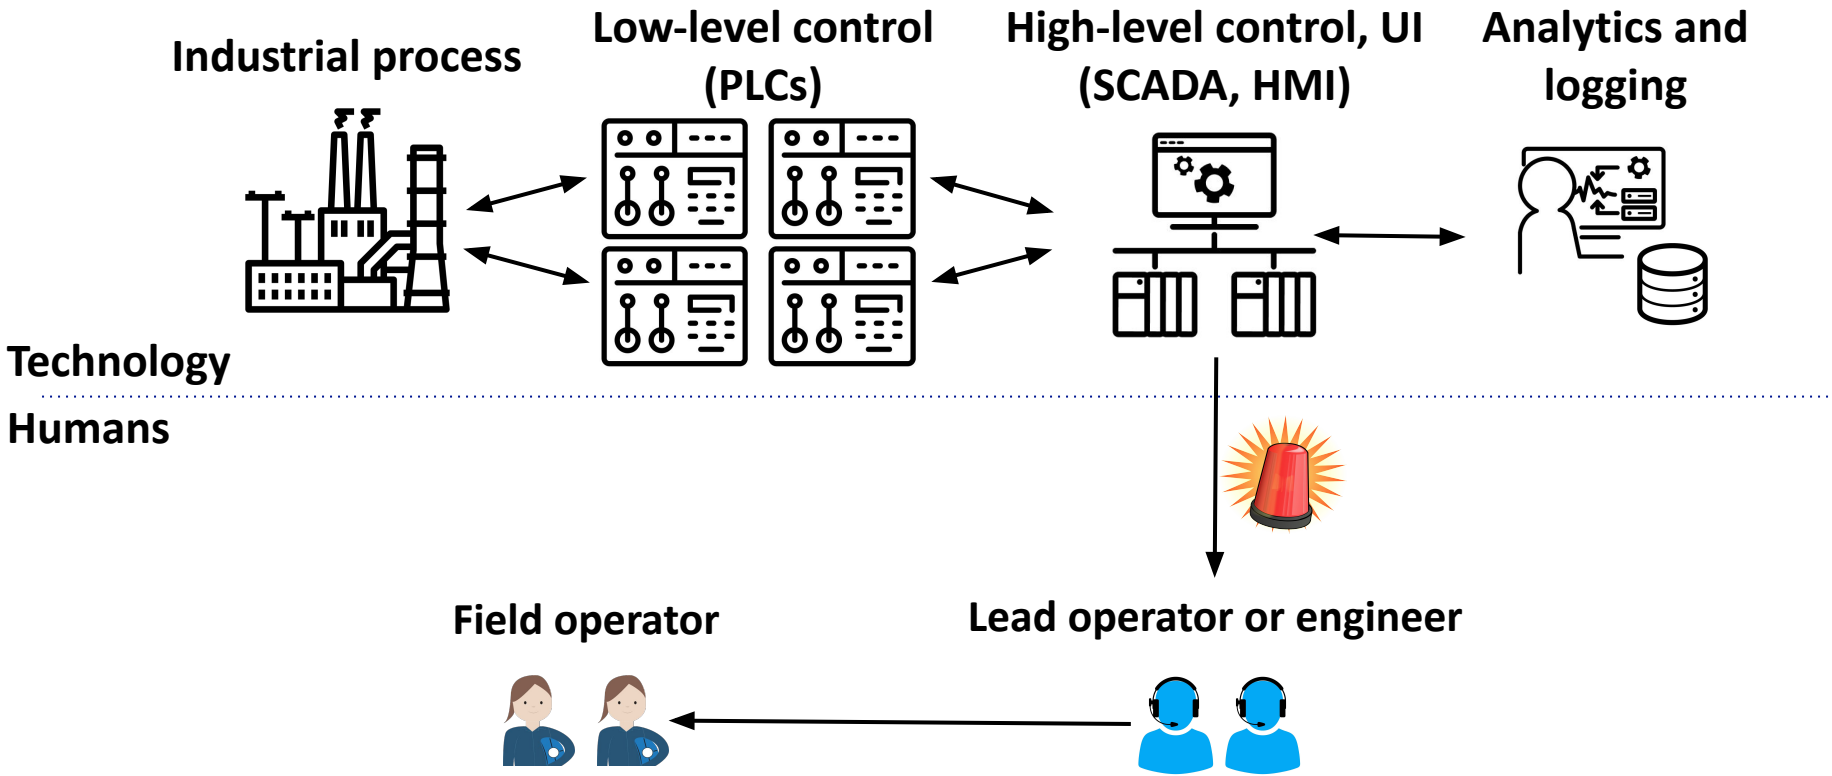

Who is involved with protecting ICS?

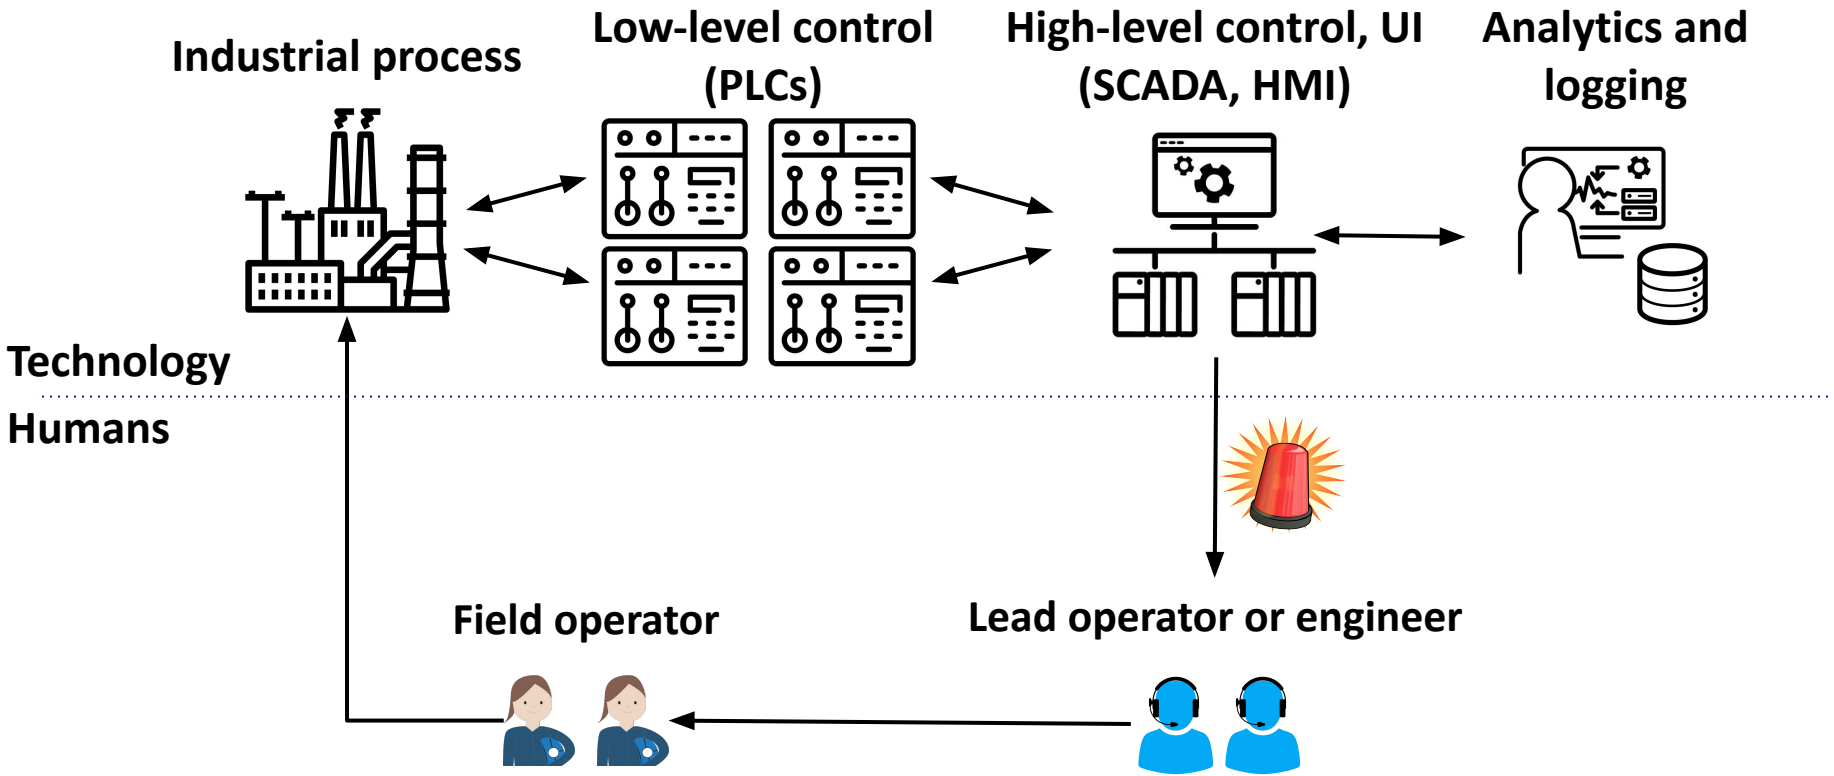

Who is involved with protecting ICS?

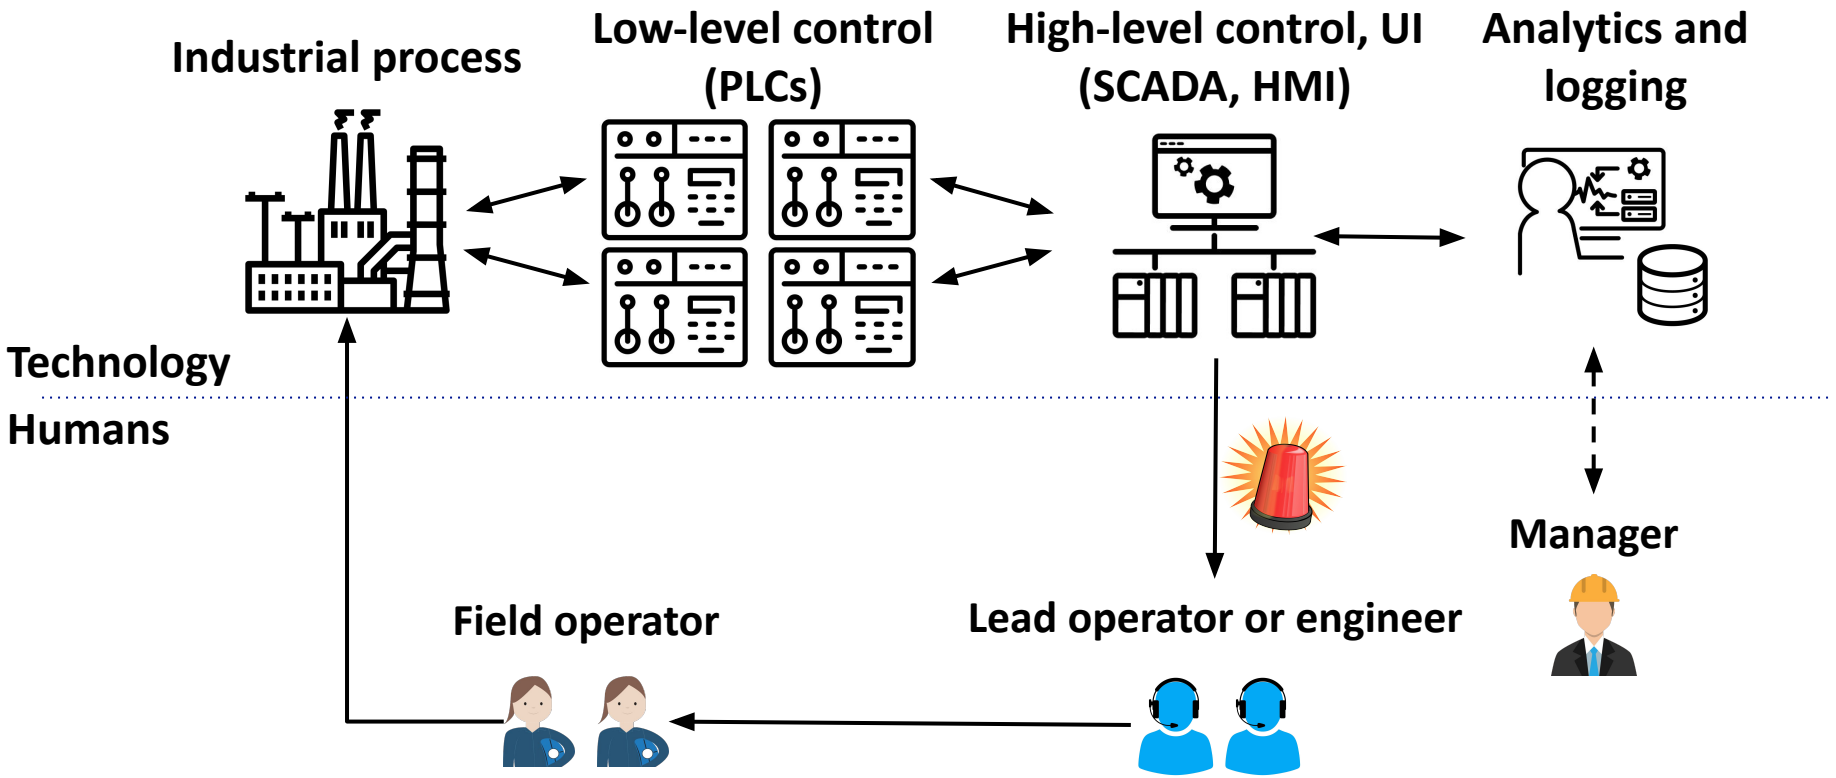

Who is involved with protecting ICS?

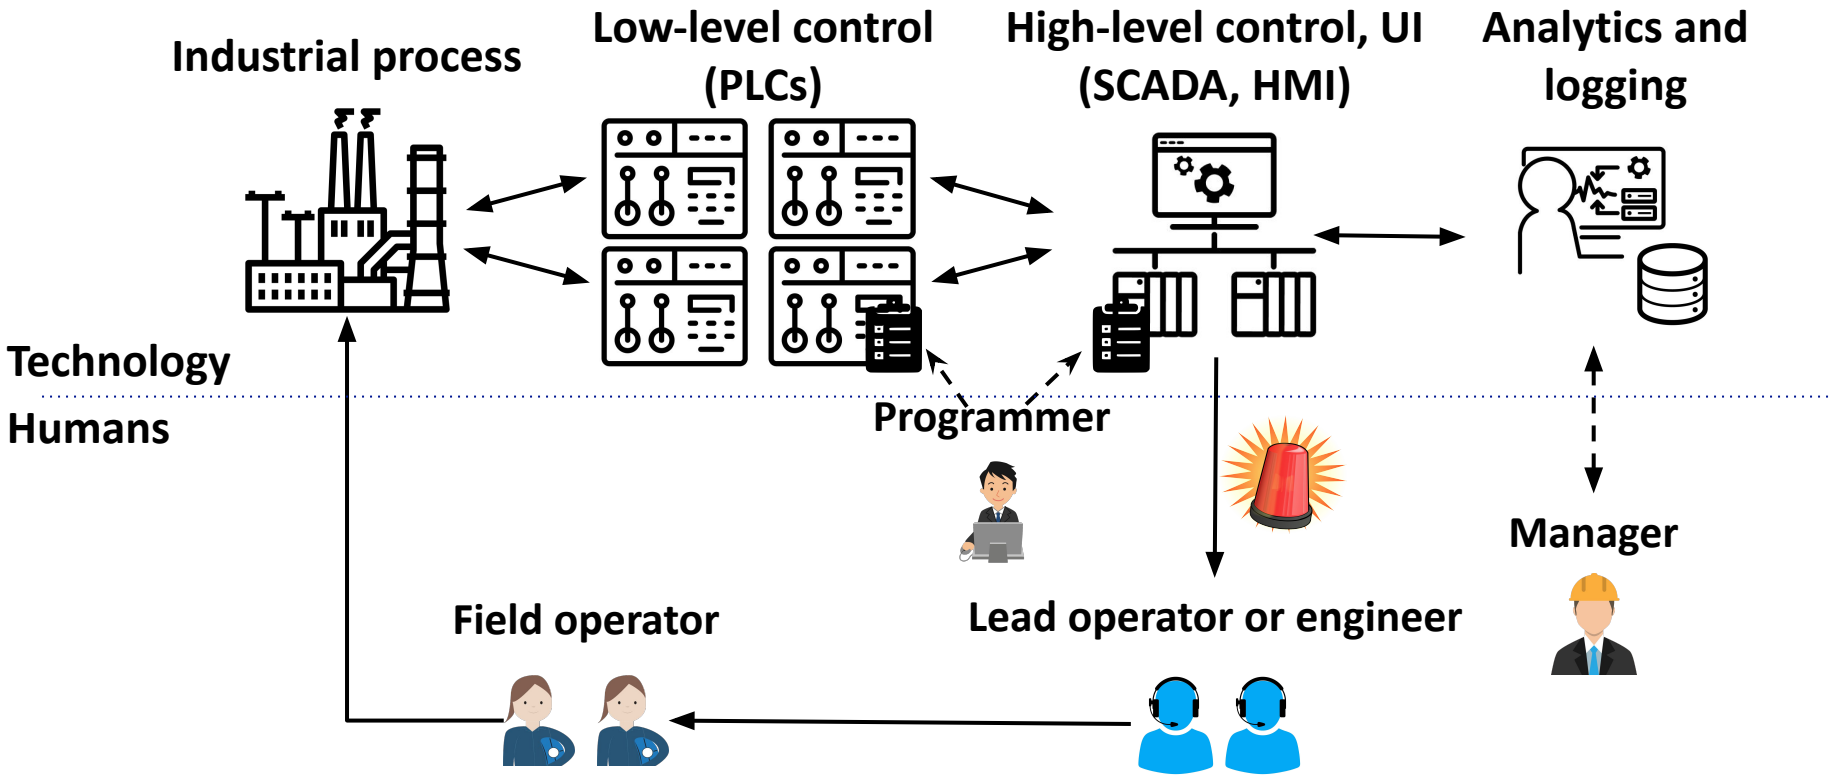

Who is involved with protecting ICS?

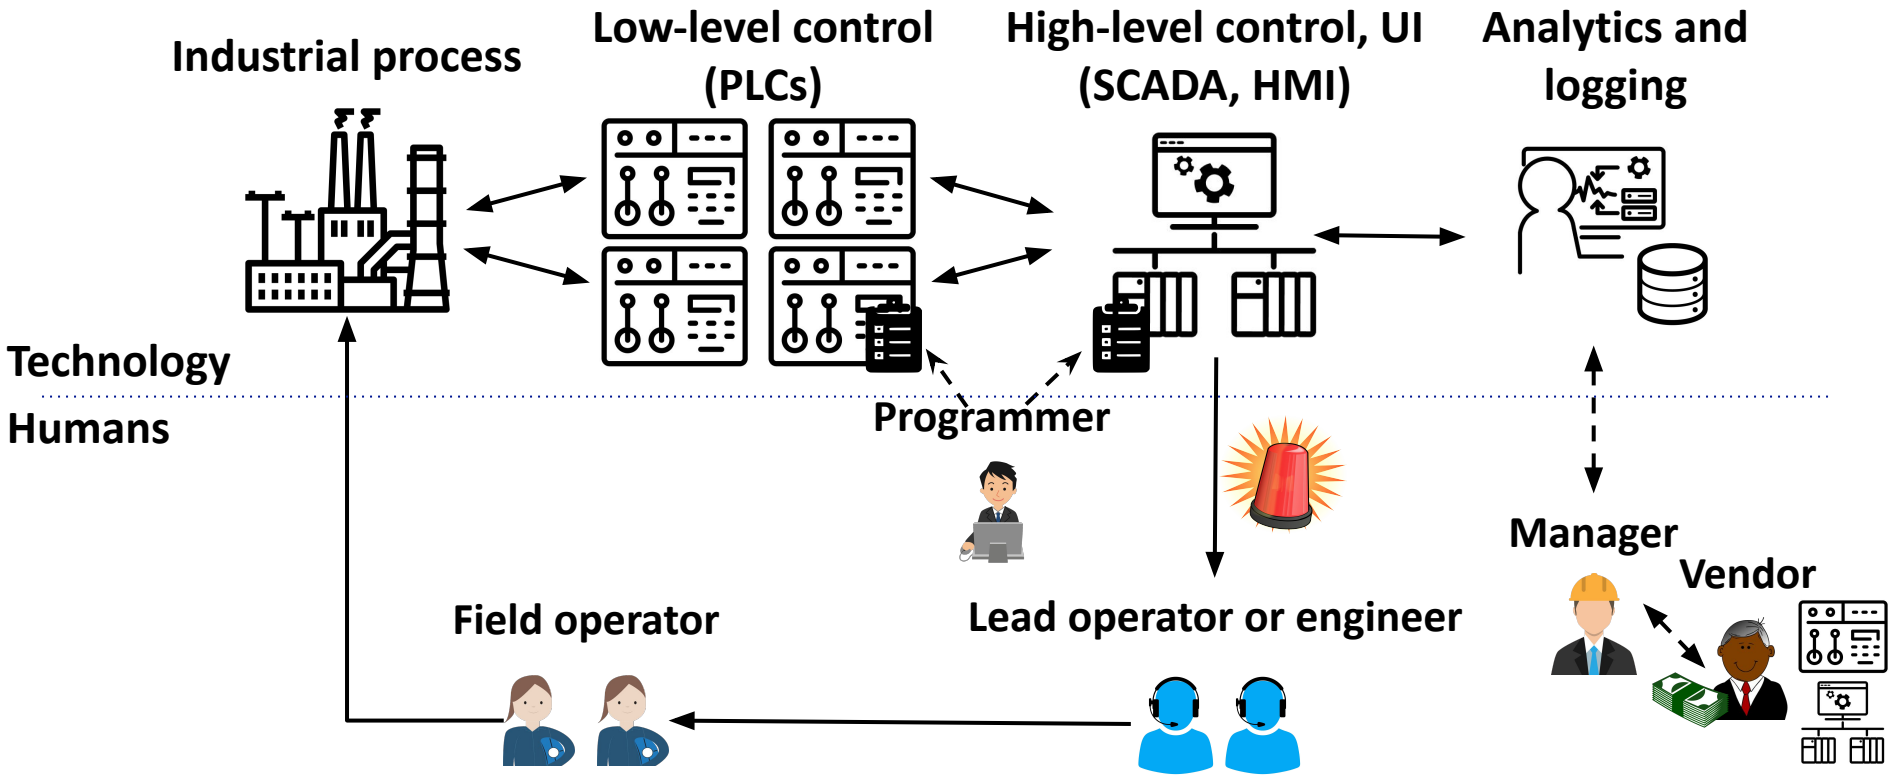

Who is involved with protecting ICS?

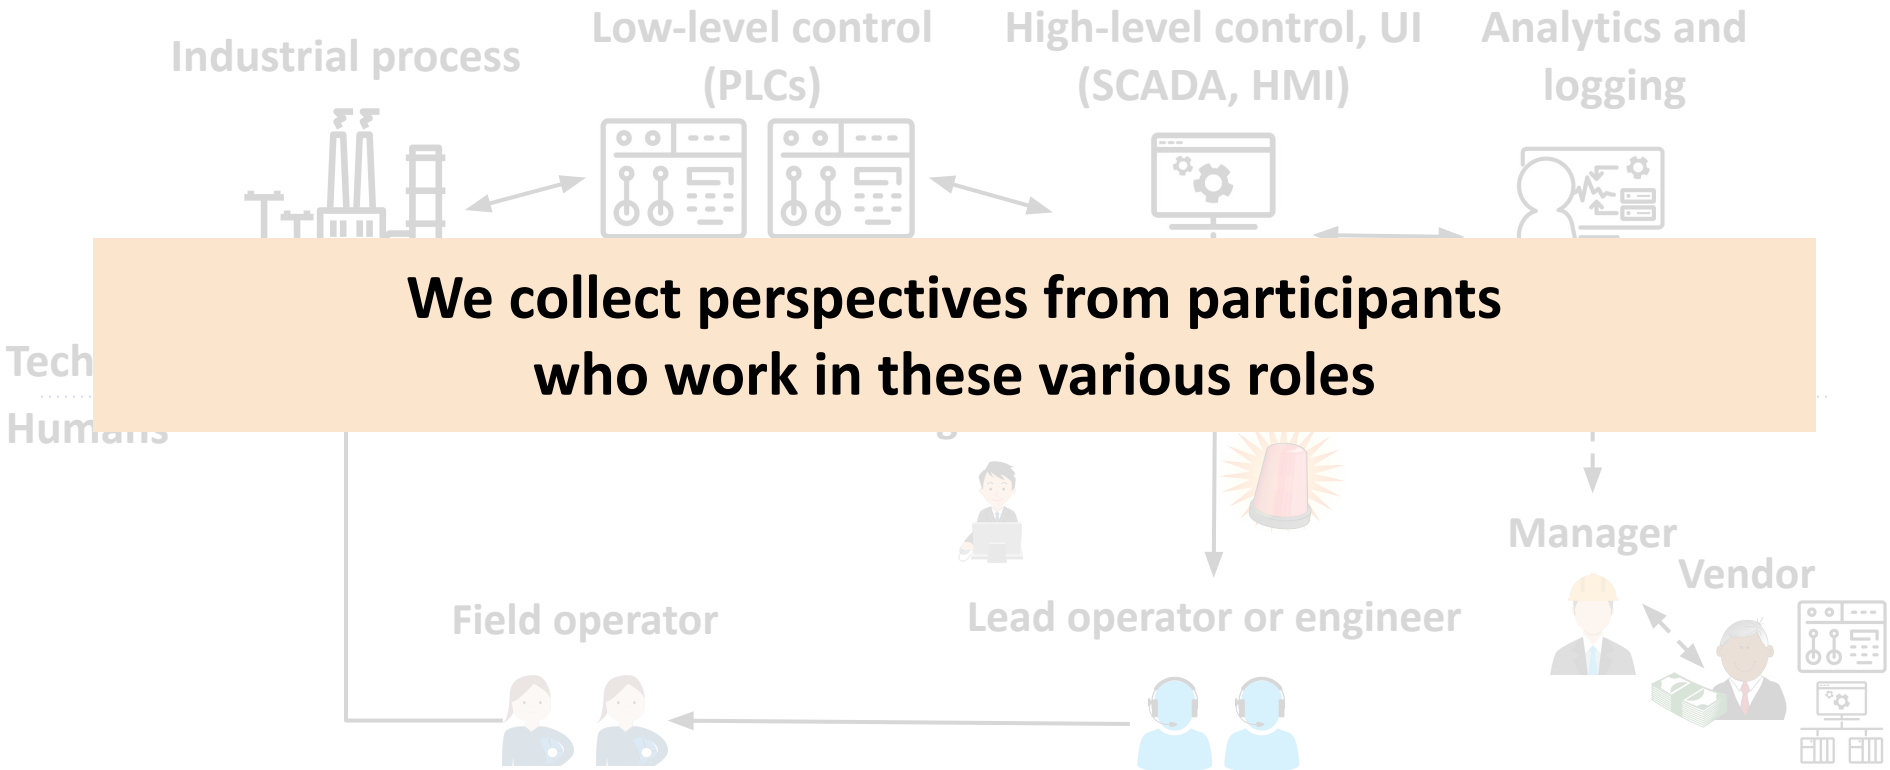

We collect perspectives from participants who work in these various roles

We recruited a variety of participants

Recruitment: 

We recruited a variety of participants

Recruitment:

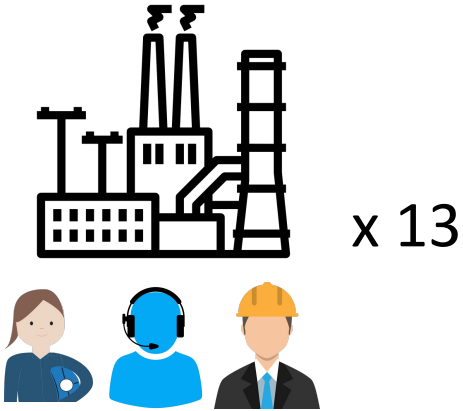

x 13

Operations

We recruited a variety of participants

Recruitment:

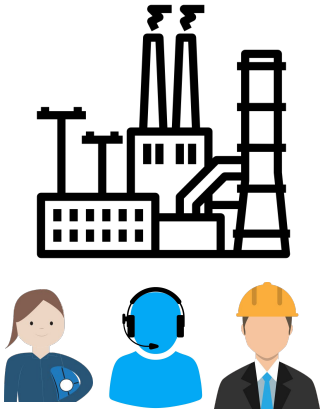

x 13

Operations

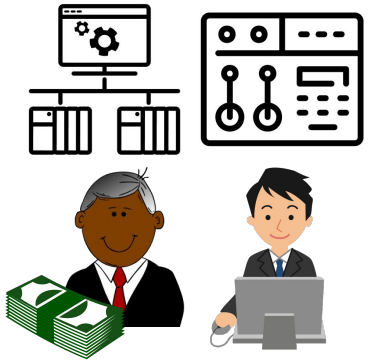

x 5

Vendors

We recruited a variety of participants

Recruitment:

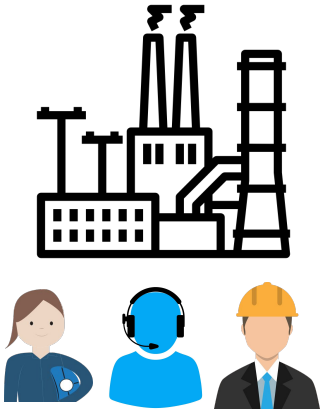

x 13

Operations

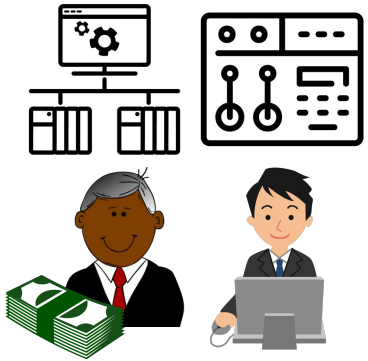

x 5

Vendors

- Electricity
- Oil and gas
- Manufacturing
- Water treatment

We use an indirect approach in interviews

- Challenge: Few practitioners with AI expertise

We use an indirect approach in interviews

- Challenge: Few practitioners with AI expertise

- **An indirect approach** via semi-structured interviews:
  - Technology and tasks for alarms
  - Adopting new technology in ICS
  - Perceptions of AI for ICS

We use an indirect approach in interviews

- Challenge: Few practitioners with AI expertise

- **An indirect approach** via semi-structured interviews:
  - Technology and tasks for alarms
  - Adopting new technology in ICS
  - Perceptions of AI for ICS

- Qualitative coding and analysis for themes relevant to AI adoption

What types of things did we learn about ICS?

What types of things did we learn about ICS?

Technology and infrastructure	• For collecting and using process data • For building alarm systems

ICS alarm systems: different shapes and sizes!

ICS alarm systems: different shapes and sizes!

ICS alarm systems use rule sets, but *how* they use them varies:

ICS alarm systems: different shapes and sizes!

ICS alarm systems use rule sets, but *how* they use them varies:

- Logic: High or low values, rate of change, or combinations

ICS alarm systems: different shapes and sizes!

ICS alarm systems use rule sets, but *how* they use them varies:

- Logic: High or low values, rate of change, or combinations
- Locations: Logic in sensors, PLCs, or SCADA
  - Forwarded and displayed in different locations

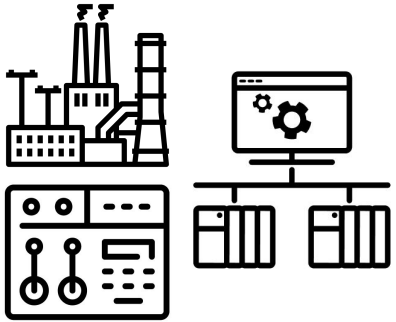

ICS alarm systems: different shapes and sizes!

ICS alarm systems use rule sets, but *how* they use them varies:

- Logic: High or low values, rate of change, or combinations
- Locations: Logic in sensors, PLCs, or SCADA
  - Forwarded and displayed in different locations
- People: Written and managed by plant owners or vendors

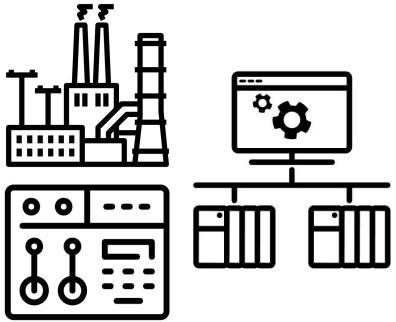

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

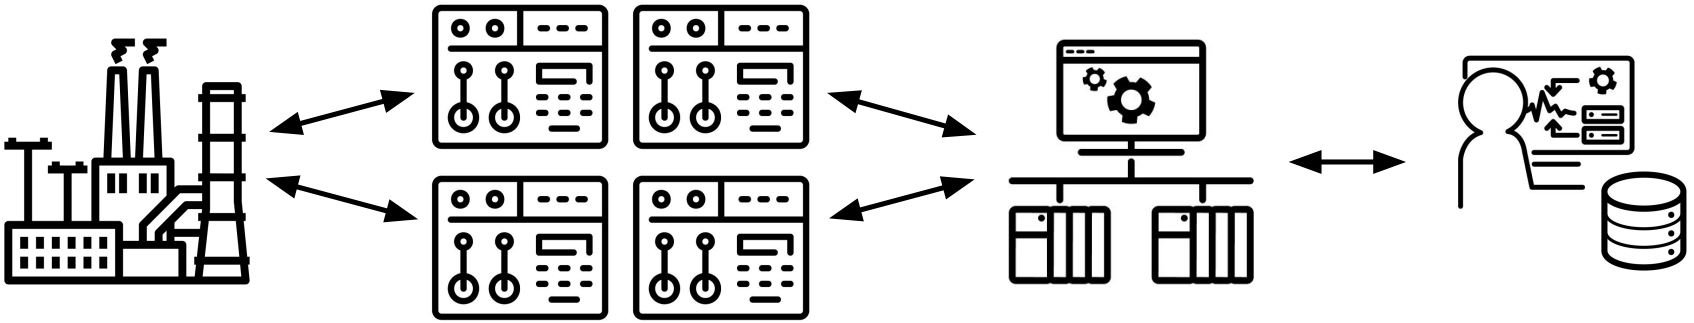

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

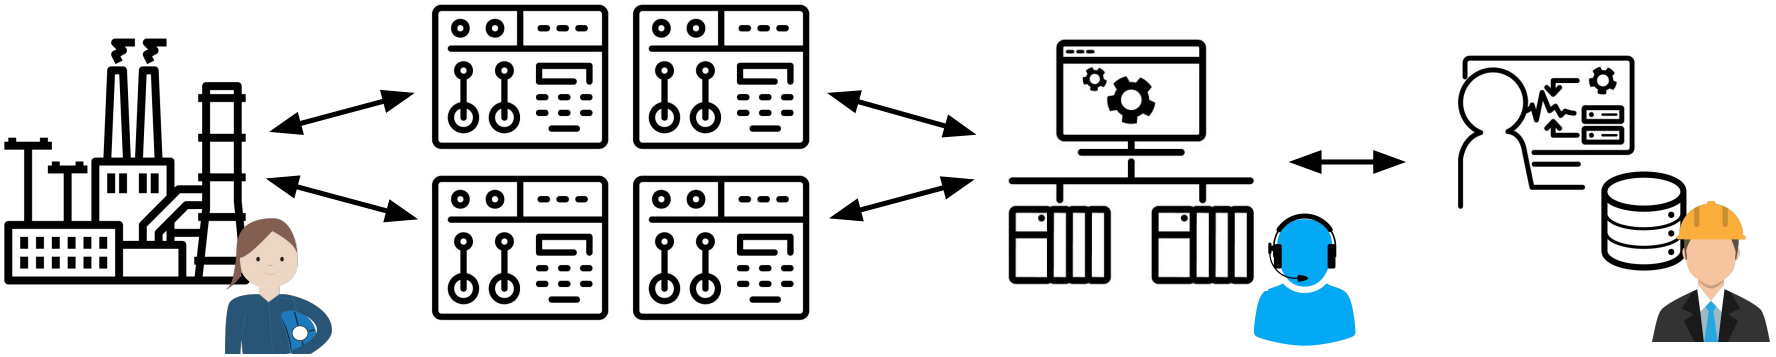

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

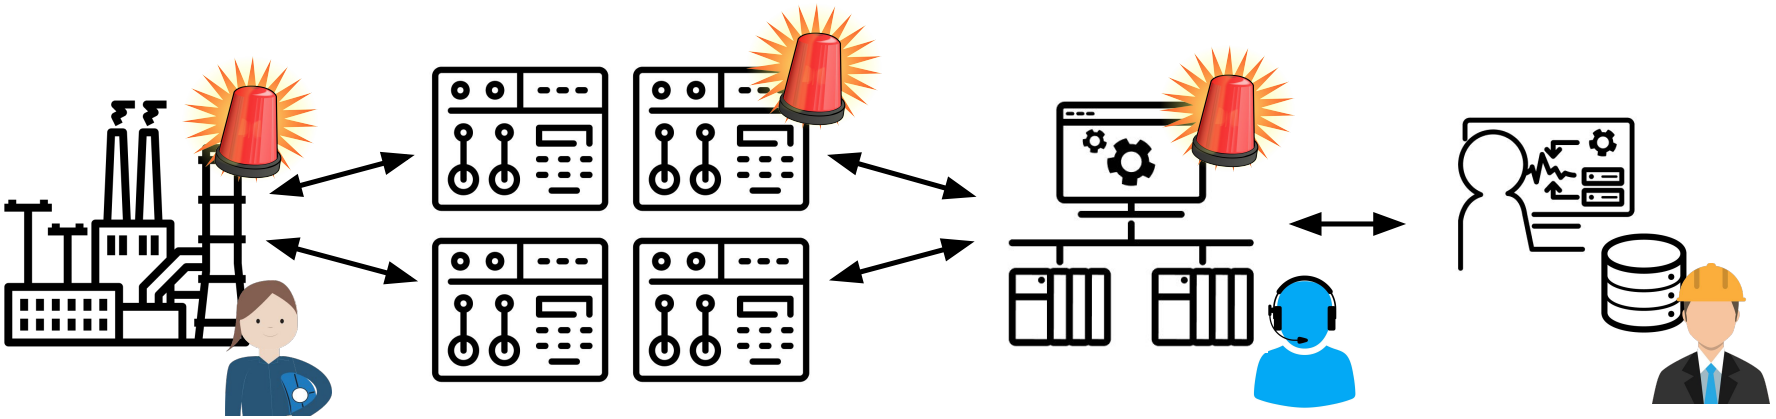

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

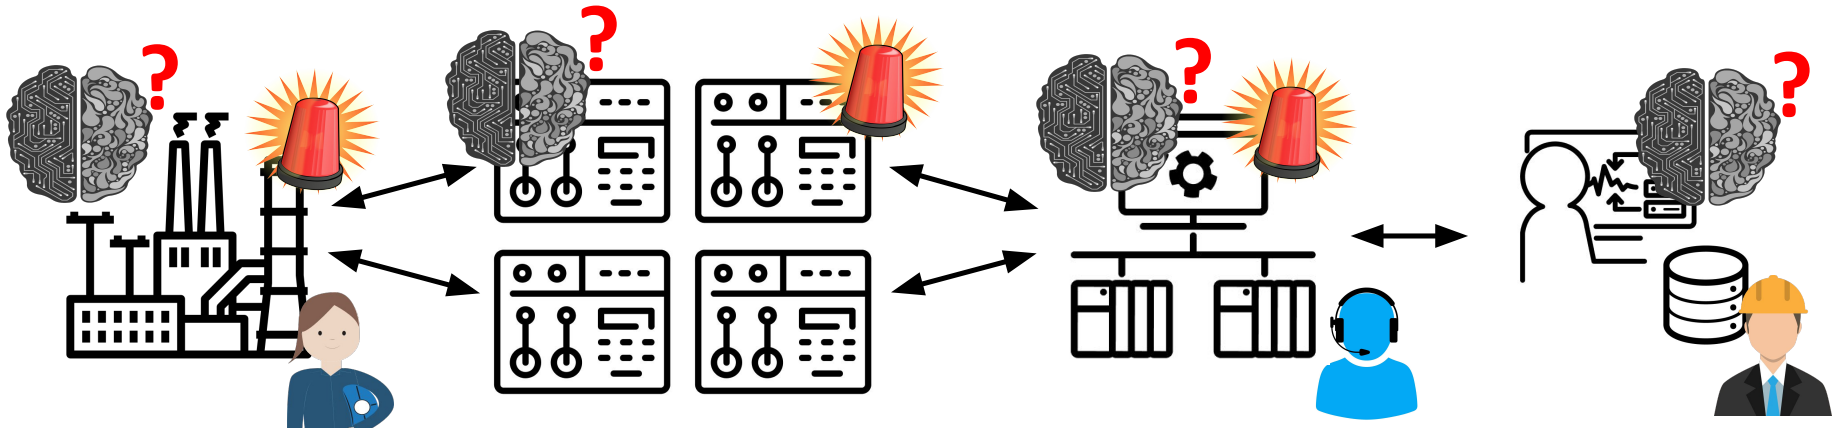

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces
- Technological and regulatory constraints

Takeaway 1: ICS setups vary, limiting the feasibility of general-purpose research solutions

Research

- *Centralized* data and compute

Practice

- *Decentralized* data, devices, and user interfaces

We need new deployment models for AI in ICS based on devices, data, and users

What types of things did we learn about ICS?

Technology and infrastructure	• For collecting and using process data • For building alarm systems
Human factors	• Human tasks involved in alarm systems • Pain points when using alarm systems • Pain points from working in ICS environments

Common human tasks in alarm workflows

Common human tasks in alarm workflows

Set up detection system

Common human tasks in alarm workflows

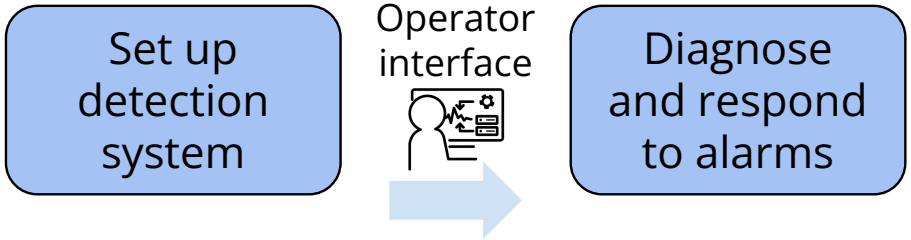

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions

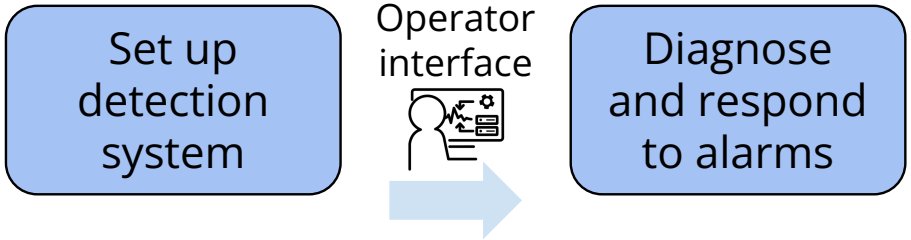

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions

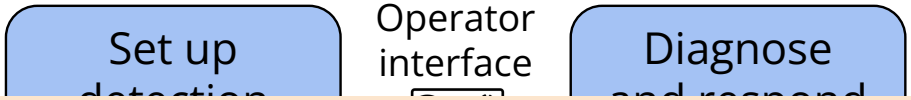

*[...] our greatest challenge is **training the staff that's still fairly new** [...] what the **appropriate level of response** is. –P18*

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions
  - Relies on **intuition and experience**

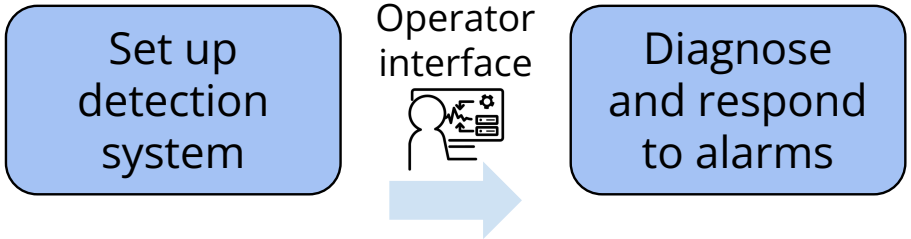

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions
  - Relies on **intuition and experience**

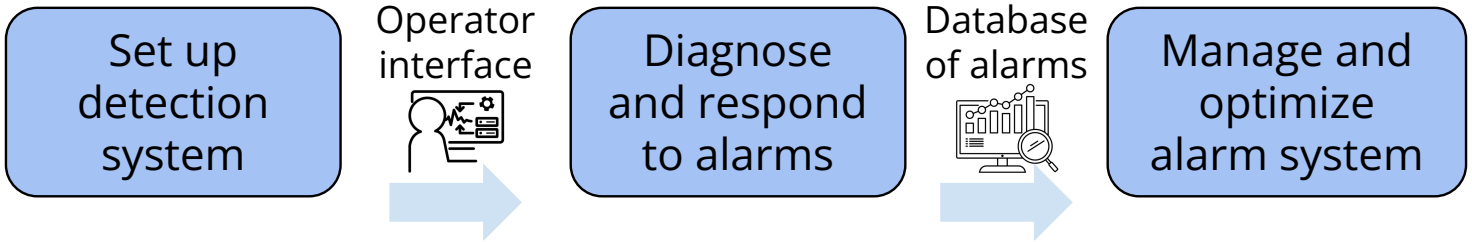

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions
  - Relies on **intuition and experience**
- Alarm management: Using prior alarm data to optimize alarm systems

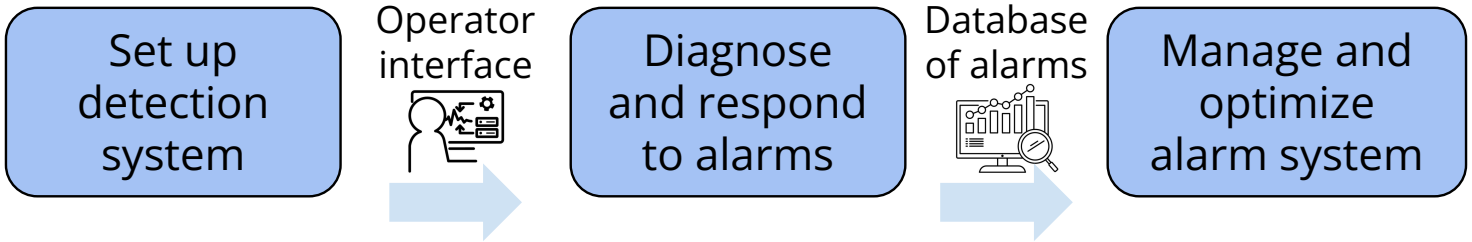

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions
  - Relies on **intuition and experience**
- Alarm management: Using prior alarm data to optimize alarm systems

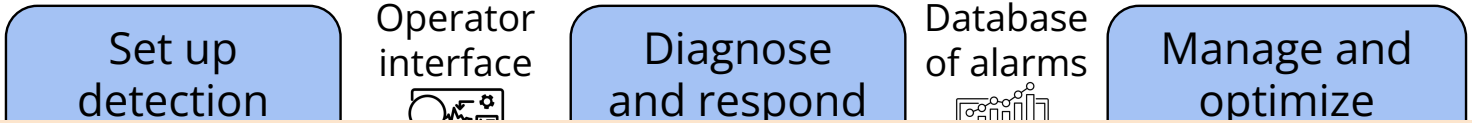

*We **looked at every single alarm** that we have, and challenged if you need the alarm, and then what the alarm point should be. And that was a significant **year and a half of, at least 10 hours a week**. –P13*

Common human tasks in alarm workflows

- Alarm diagnosis: Determining and performing follow-up actions
  - Relies on **intuition and experience**
- Alarm management: Using prior alarm data to optimize alarm systems
  - Determining **what should be an alarm** is difficult

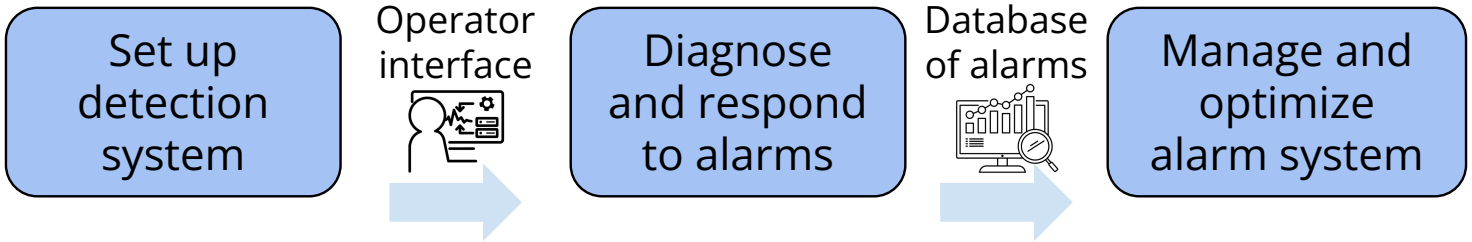

Takeaway 2: Operators want help with tasks beyond anomaly detection

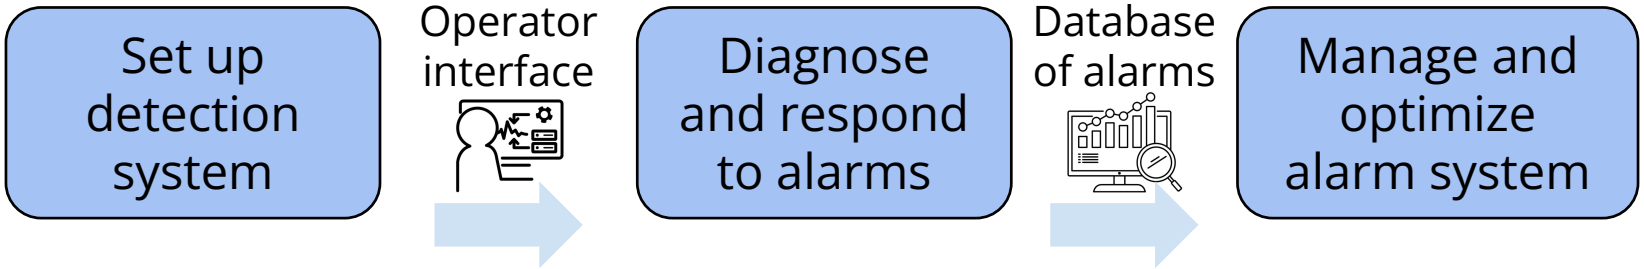

Takeaway 2: Operators want help with tasks beyond anomaly detection

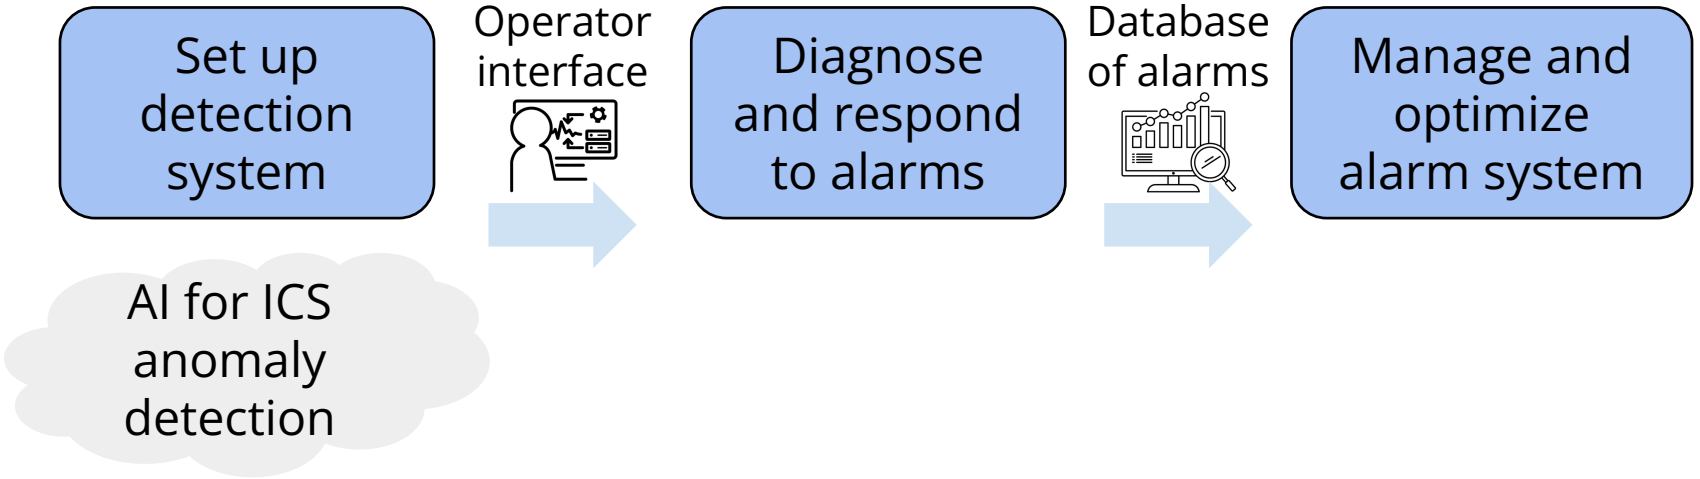

Takeaway 2: Operators want help with tasks beyond anomaly detection

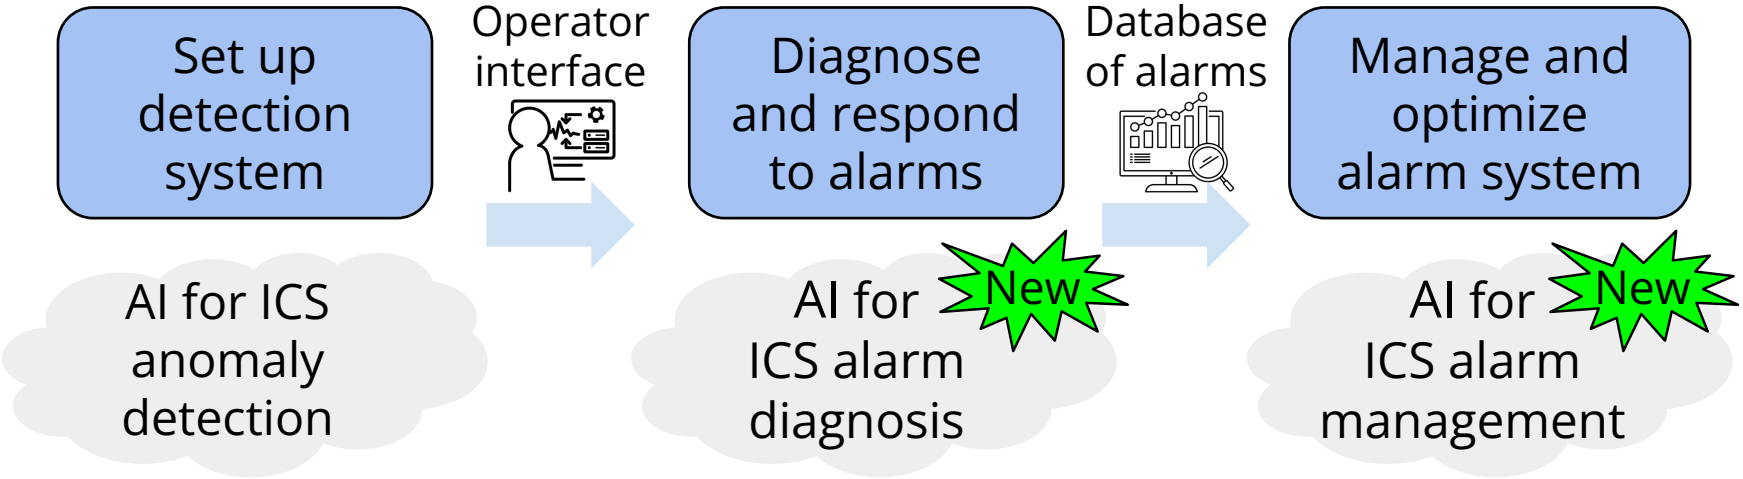

What types of things did we learn about ICS?

Technology and infrastructure	• For collecting and using process data • For building alarm systems
Human factors	• Human tasks involved in alarm systems • Pain points when using alarm systems • Pain points from working in ICS environments
AI adoption	• Perspectives on vendors and tool adoption • Perspectives on how AI could help them

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible
  - For detection, as proposed in research

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible
  - ~~For detection, as proposed in research~~
  - For non-critical, complex tasks
  - Use AI to assist and make suggestions

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible
  - ~~For detection, as proposed in research~~
  - For non-critical, complex tasks
  - Use AI to assist and make suggestions

We now have 1000s of examples of: the data was doing this, it led to this root cause analysis, and it led to this action. [...] we can begin to look at applying deep learning, because ***we have the necessary data to train that****. –P1*

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible
  - ~~For detection, as proposed in research~~
  - For non-critical, complex tasks
  - Use AI to assist and make suggestions

- Almost all practitioners have reservations about AI
  - Overconfident, inaccurate, difficult to understand

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Some belief that adopting AI is feasible
  - ~~For detection, as proposed in research~~
  - For non-critical, complex tasks
  - Use AI to assist and make suggestions

- Almost all practitioners have reservations about AI
  - Overconfident, inaccurate, difficult to understand

The management of our plant, they don't really trust AI because they ***don't have a solid understanding of how it works****. –P17*

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Most common request about AI: more transparency about how AI works

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Most common request about AI: more transparency about how AI works

*I need to be able to get in there and **do some development or make changes** and what that looks like is going to make me a lot more comfortable. –P4*

Takeaway 3: Practitioners are optimistic about AI's potential, if introduced carefully

- Most common request about AI: more transparency about how AI works

*I need to be able to get in there and **do some development or make changes** and what that looks like is going to make me a lot more comfortable. –P4*

*To implement [AI], the good ways all involve: Here's how it works, here's what it's looking at, breaking it down and **putting a lot more transparency behind it**. –P11*

Adopting AI to Protect Industrial Control Systems:
Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

Adopting AI to Protect Industrial Control Systems:
Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) **We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption**

Adopting AI to Protect Industrial Control Systems:
Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) **We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption**

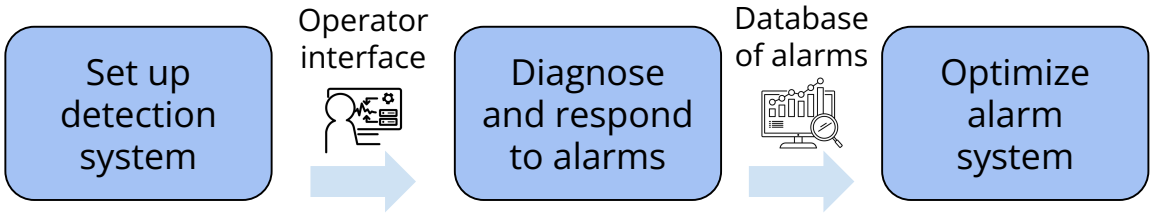

Adopting AI to Protect Industrial Control Systems:
Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) **We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption**

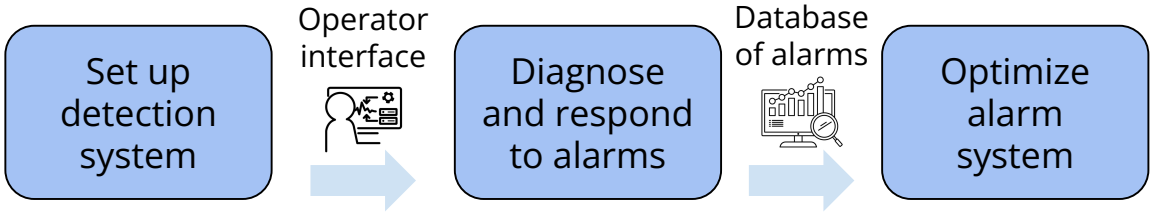

2) **We recommend that researchers in AI for ICS:**

Adopting AI to Protect Industrial Control Systems: Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) **We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption**

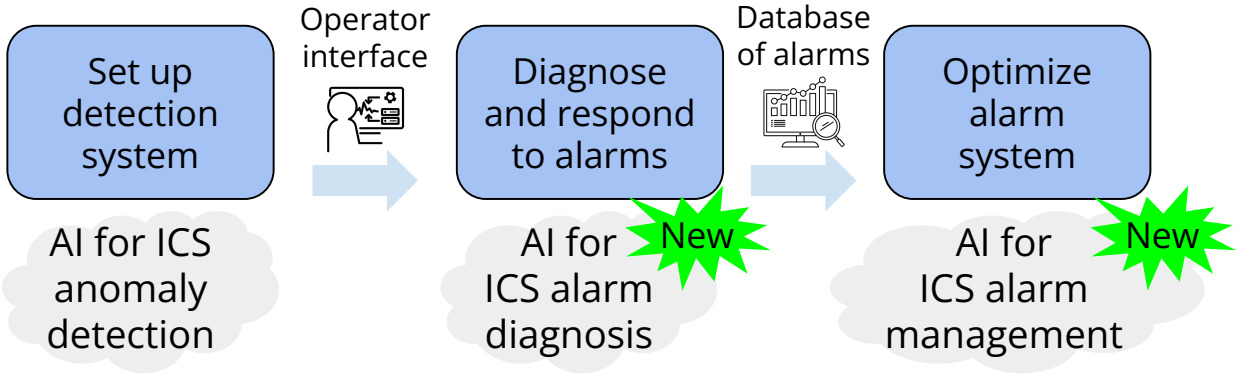

2) **We recommend that researchers in AI for ICS:**

- Focus on alarm diagnosis and management

Adopting AI to Protect Industrial Control Systems:
Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption

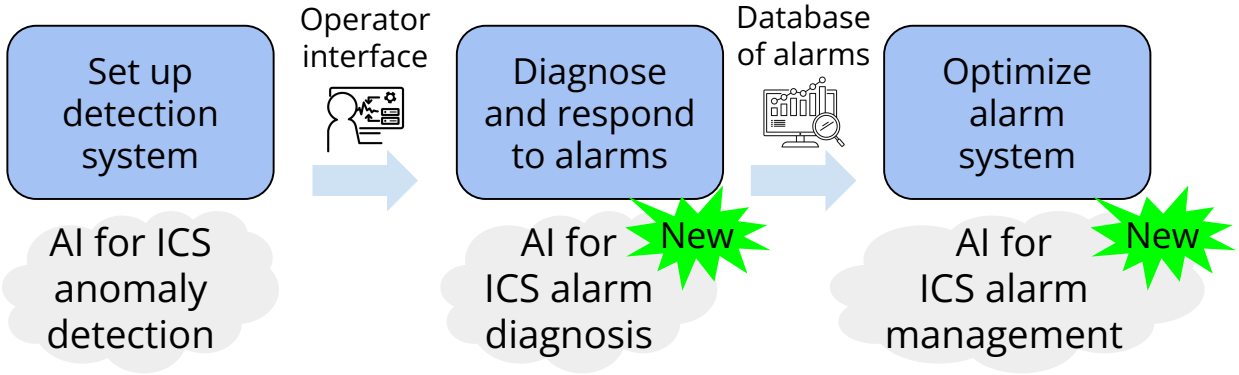

2) We recommend that researchers in AI for ICS:

- Focus on alarm diagnosis and management
- Consider technical and regulatory constraints on data collection

Adopting AI to Protect Industrial Control Systems: Assessing Challenges and Opportunities from the Operators' Perspective

Clement Fung, Eric Zeng, Lujo Bauer
Contact: https://clementfung.me

1) We identify differences between AI research and ICS practice—technical and human factors—that hinder AI adoption

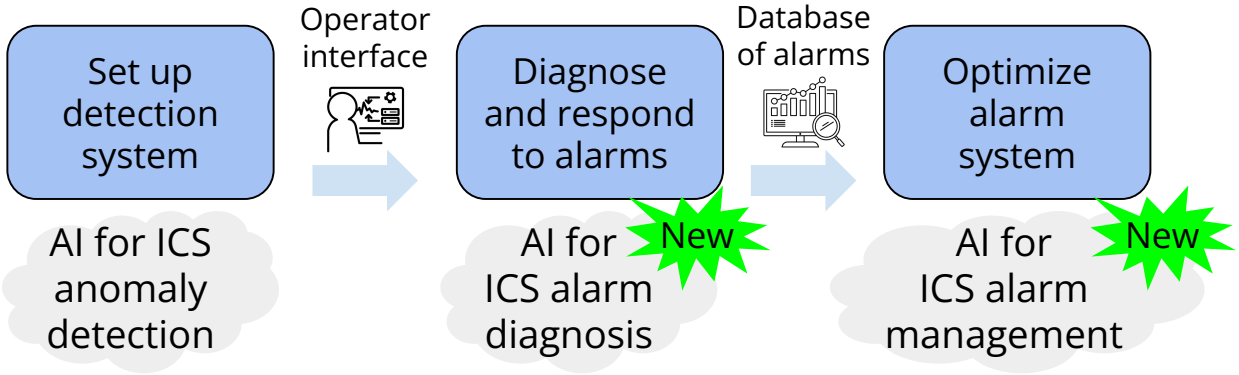

2) We recommend that researchers in AI for ICS:

- Focus on alarm diagnosis and management
- Consider technical and regulatory constraints on data collection
- Demonstrate AI transparency through interactive pilot projects