Securing Industrial Systems with the Integration of Neuro-Symbolic AI

Ghena Barakat
Supervisors: Prof. Antonio Puliafito & Prof. Giovanni Merlino

Executive Summary

- Industrial Control Systems (ICS) and Operational Technology (OT) are critical infrastructures
- Cyberattacks on industrial systems are increasing in frequency and sophistication
- This research proposes a Neuro-Symbolic AI approach for industrial cybersecurity
- The approach combines:
  - Learning-based anomaly detection
  - Symbolic reasoning based on industrial domain knowledge

Industrial Cybersecurity Landscape

- Industrial systems differ from IT systems due to:
  - Physical processes and safety constraints
  - Real-time operation requirements
  - Long equipment lifecycles
- Common attack targets:
  - Industrial networks
  - Sensors and actuators
  - PLCs and control logic

Problem Statement

Traditional cybersecurity solutions struggle in industrial environments due to:

- Deterministic industrial protocols (Modbus, DNP3, OPC-UA)
- Limited labeled attack data
- Requirement for explainable decisions in safety-critical systems
- Strict real-time constraints and low false-positive tolerance
- Presence of air-gapped and legacy systems

Why Pure AI or Pure Rules Are Not Enough

- Pure Machine Learning:
  - Requires large labeled datasets
  - Often black-box and hard to trust
- Rule-Based Systems:
  - Interpretable but rigid
  - Difficult to adapt to new attack patterns
- Industrial cybersecurity needs:
  - Learning + Reasoning
  - Adaptability + Interpretability

Proposed Solution: Neuro-Symbolic AI

A hybrid cybersecurity architecture combining:

- Neural networks for learning patterns and anomalies
- Symbolic reasoning for validation and explanation
- Tight integration between learning and logic

*Goal: Accurate, explainable, and reliable intrusion detection in industrial systems

Neuro-Symbolic AI Architecture (Overview)

- The proposed system follows a Neuro-Symbolic architecture that integrates learning and reasoning
- Learning component:
  - Learns normal system behavior from operational data
  - Identifies deviations that may indicate security threats
- Reasoning component:
  - Uses formal rules and domain knowledge
  - Ensures decisions are consistent with system logic and constraints

Neuro-Symbolic AI Architecture (Overview)

- Integrated decision process:
  - Combines learned patterns with logical validation
  - Produces reliable and explainable security alerts

*Outcome: A balanced approach that improves detection accuracy while maintaining transparency and trust.

Future Work: Extending to Federated Learning

- Industrial data is sensitive and often siloed
- Federated Learning enables:
  - Collaborative model training across sites
  - No raw data sharing
- Benefits:
  - Improved generalization
  - Privacy-preserving security intelligence sharing
- Neuro-symbolic models can be updated collaboratively across factories

Future Work: Adapting to Tiny Models

- Many industrial devices are resource-constrained
- Tiny AI enables:
  - On-device inference
  - Low-latency security monitoring

*Goal: Lightweight neuro-symbolic models for edge deployment

Conclusion

- Industrial cybersecurity requires more than black-box AI
- Neuro-symbolic AI provides:
  - Learning
  - Reasoning
  - Explainability
- Future integration with:
  - Federated Learning
  - Tiny Models
- A promising path toward secure and trustworthy industrial systems