Lessons Learned for AI in Plant OT and IACS

Industrial control and optimization is projected to be among the most highly profitable applications of “real world” Artificial Intelligence (AI). However, as “AI agents” gain direct or even indirect control of such critical infrastructure, the risks of a malfunction or malicious acts rise even more quickly than potential benefits. To address these risks, corporate leaders of this emerging technology have urged that governments take measures to ensure the safety of AI, especially where it controls critical infrastructure.

Several of these AI leaders have even suggested that government regulation is required before any such applications are implemented, and that a “kill switch” should be a minimum requirement of such AI systems.

While a public dialog on these topics is necessary, I believe that government regulation is both impractical and unnecessary.

- **It is impractical** because such regulations typically take years to formulate and AI is developing so quickly that regulations cannot possibly get ahead of AI implementation. Even worse, if the current world AI leaders in the US and Europe “hold back”, they will certainly be overtaken by less capable but more authoritarian regimes increasing the risks instead of reducing them.

- **It is unnecessary** because industry has been applying increasingly sophisticated control and optimization applications for many decades. In fact, it can be argued that Artificial Intelligence (AI) is not a revolution, but rather an evolution of increasingly sophisticated matrix mathematical solutions that have been made possible by ever more powerful computers.

Rather than talk in generalities, I thought it would be better to provide actual examples from large-scale Operational Technology (OT) and Industrial Automation and Control Systems (IACS) that I have led. In each case, although these OT and IACS introduced risks and uncertainty, they were effectively dealt with by engineers with the necessary industrial experience. The real risk is not AI, but rather that it may be applied without learning from the lessons of the last 50 years.

Here are a few examples:

- In the early 1970's, **Linear Programming Optimizers** were used for optimizing shipping schedules and refinery production planning. These used large linear matrices to maximize an objective function (usually profit in $). These models were so large that they would "take over" our entire corporate data center (then an IBM 370-158) for nearly an hour.
    - **Issue:** Local maxima and simple data errors could have large economic impacts.
    - **Resolution:** Humans who had done these schedules in the past, reviewed recommended actions and deleted "outliers".

- In the mid 1970's **Non-Linear optimizers** allowed even more sophisticated relationships to be modelled and maximized for problems such as gasoline blending.
    - **Issue**: Since relationships could be non-linear, there was no practical way for even experienced blenders to reject invalid optima.
    - **Resolution**: "Test engines" were used to verify proposed blends in a routine quality control procedure before actual blending of large volumes of product.

- By 1975 **advanced control algorithms** were able to stabilize fast loops and analyzer feedback more quickly and efficiently than traditional PID control.
    - **Issue**: If advanced controllers fail, the operator must quickly take over "manual control". In an early computerized Peroxide Reactor a programmer failed to provide a "manual" mode for the operator, and the resulting explosion took the entire polyethylene plant off line for a month.
    - **Resolution**: Virtually all controllers provide "Auto", "Manual" and "Off" switches, and where necessary, High and Low Alarms and automatic safety interlocks as backup. No competent control engineer would depend on a "Kill switch" (as has been proposed by senior AI leaders to keep AI systems "safe").

- By 1980 **advanced digital control** was developed that could stabilize interacting loops and incorporate analyzer feedback more quickly and efficiently than operators and traditional PID control.
    - **Issue**: If model-based control failed, the operator of a high pressure polyethylene reactor could not control pressure quickly enough to prevent a "decomp" that was known to break windows a mile from the plant.
    - **Resolution**: On loss of reactor control, automated high speed "sequences" safely "ramp down" the unit. These sequences

recognize the plant status and automatically, ramp, hold, resume or restart the high pressure circuit. Activation of an "AI kill switch" would very likely destroy the ethylene compressor building killing everyone in it.

- In the 1990s **Expert Systems** used a large number of "rules" provided by the "best" operator(s) to control their process units. A large matrix algorithm then used real-time process conditions as inputs to find a "best solution" for process performance.
    - **Issue:** Operators often did not understand the resulting plant control actions.
    - **Resolution:** Traditional "alarms" and "alarm procedure books" ensured that unsafe situations were dealt with promptly. As a further backup, the "Lead Operator" supervised actions of the Expert System as he would have done for the replaced operator.

- In the 2010's **"digital twin" simulations** of the plant were available that could anticipate operating problems and "pre-test" proposed solutions to plant upsets.
    - **Issue**: Digital Twin simulations were more sophisticated than operators could quickly understand. Even plant engineers often could not modify or troubleshoot these "digital twin" models.
    - **Resolution**: High speed networking technologies allowed Remote Process specialists to quickly provide support. Emergency response centers (ERC) are often also responsible for development and maintenance of the digital twin model.

- In 2020 we began to see **cyber attacks** on industrial infrastructure that were faster and probed more vulnerabilities than a human might accomplish. In order to cope with these, AI-based cyber defense systems will increasingly be required.
    - **Issue**: Plant Operators will not be able to respond quickly enough to participate in cyber defense before it is over. In fact, they may never even realize that an attacker had penetrated their facility.
    - **Resolution:** Design of control systems must incorporate measures to deal with both malicious and inadvertent events. Network Operations Centers (NOCs) will be needed to watch for unusual traffic and respond when cyber incidents are suspected.

It is important to realize that in each of the above examples, safe operation requires a deep understanding of the hazards inherent in plant facilities, applications and equipment. What matters is not the optimization technology used. Instead, what is required, is a clear understanding of, state-based and sequence based control, human interface design, and the practices and standards of each industry.

AI applications are now gaining the ability to “self-improve” without human direction. This is potentially both a blessing and a curse. It may be beneficial, making AI more effective in industrial applications; however, it also brings the risk that AI applications may “improve” plant operations in ways that were not intended. We must therefore find ways for humans to retain the necessary information and ability to intervene.

Articles that address the use of AI in IACS and OT, are addressed in further white papers that may be found at https://www.pera.net/Ind_tech.html

- Lessons for AI in Plant OT and IACS
- AI Plant Interface Concepts
- Neural Net Plant Optimization
- Plant monitoring, alarm and KPI reporting
- Process Operator Copilot
- AI in Emergency Response Centers
- AI Integrated Engineering
- AI Training Assistants
- Large Language Models and Knowledge Retrieval