ADVANCED REACTOR SAFEGUARDS & SECURITY

# Data-driven optimization of PPS


Presented by:

Nathan Shoman[1], Alan Evans[1], Kanad Khanna[1]

[1] Sandia National Laboratories


May 15, 2024


Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.



Contact: nshoman@sandia.gov



SAND NO. 2024-05874PE

# Facility physical protection design can be slow and expensive


- Physical security can be a significant component of operation and maintenance costs for nuclear power plants
  - Consequently, optimizing for costs while retaining effective security is an ongoing development priority
- Designing physical protection systems (PPS) can take considerable time and rely on expert judgement
- PPS design can be thought of as a large-scale optimization problem
- New approaches and tools could accelerate the development cycle and resulting in cheaper, but more effective designs

# Learn by playing: reinforcement learning

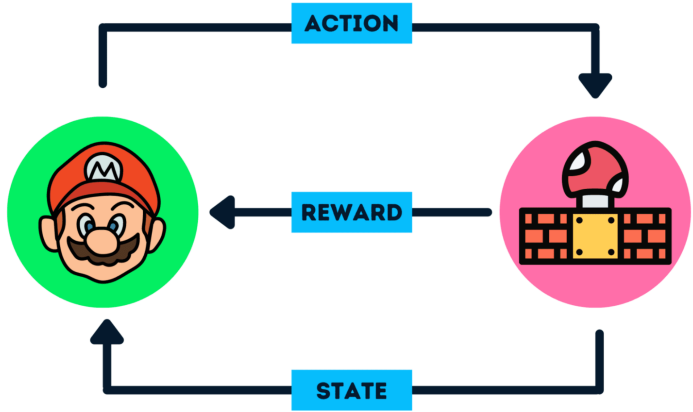


Image credit: Kdnuggets

# Agent-based exploration can help more thoroughly explore the PPS design landscape

- **Benefit:** More thoroughly explore the design space
- PPS design is a complex problem relying on expert judgement combined with trial and error (i.e., system evaluations)
- Reinforcement learning can bring a principled exploration of the design space to the problem

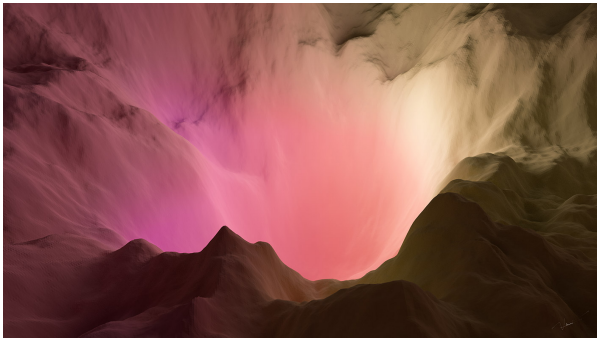

Image credit: Losslandscape

# Reinforcement learning can also help optimize for different performance targets


- Reward structure can be changed to prioritize different goals (i.e., cost, physical footprint, etc.)
- Agents can dynamically explore environment in real-time and react optimally under different conditions (e.g., before/after detection)
- Reinforcement learning can explore different states of knowledge of adversaries (e.g., knowledge of facility layout)

**Project goal**: Develop designer and adversarial agents

- Designer agent will propose candidate PPS layouts depending on criteria set by user
- Adversarial agent will find optimal paths and severe vulnerabilities

# Project roadmap


- **FY24:** Develop a baseline adversarial agent and compare to non-ML approaches (A*, RRS) and existing PathTrace algorithm (Dijkstra on directed graph).
  - Adversarial agent should produce identical results to existing algorithm demonstrating that agent is well trained and Dijkstra finds the most optimal path under prescribed constraints.
  - Requires build out of environment for agent to learn (including size optimization), policy optimization algorithm evaluation, and developing well specified reward structure
- **FY25/FY26:** Development of designer agent with comparison to existing, human-designed PPS systems.

# Fundamental challenges

- Asymmetric agent action set
- Joint training of agents
- Robust reward structure specification
- Environment size

Image credit: Deepmind

# Problem accelerators and risk reduction

- Existing state-of-the-art (PathTrace) already well suited for game playing (discretized grid, model file specification)
- Relatively easy baseline assumptions
  - Adversary has complete knowledge of facility, PPS elements, and target location
- Adversarial agents reduce the need for large quantities of human-designed PathTrace models

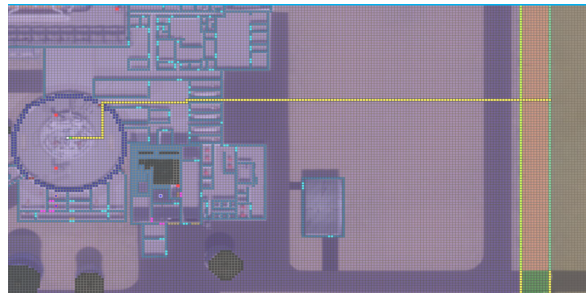

Image credit: SNL ModSim

# Progress to date (not much!)

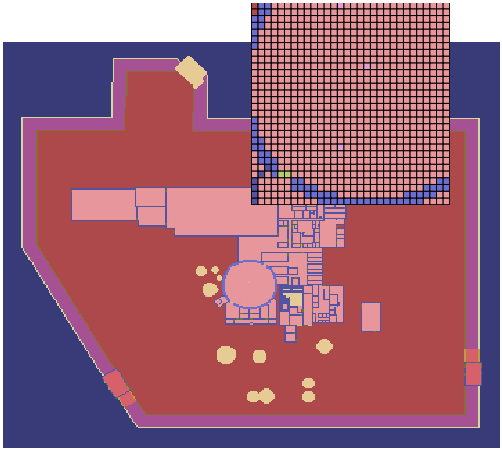

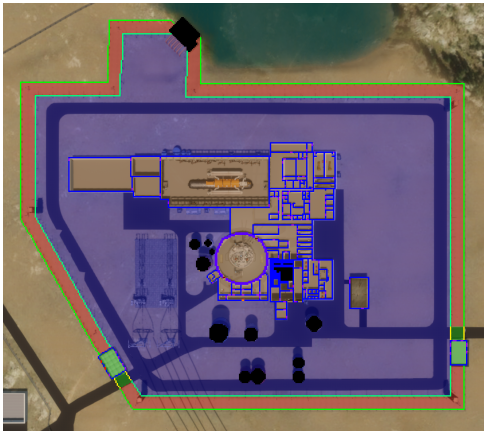

# Summary


## New project in FY24 to use ML for accelerated PPS design

- If successful, ML agents will develop candidate PPS designs using a flexible set of reward weightings
- Agents will be able to more thoroughly explore the design landscape while reacting dynamically to scenario conditions
- Design agent should be able to optimize for cost, footprint, response force response time, and more


Questions?