SAND2019-9157PE

AlphaGrid: Grid Stability Using Machine Learning State Space Navigation

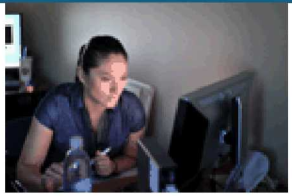

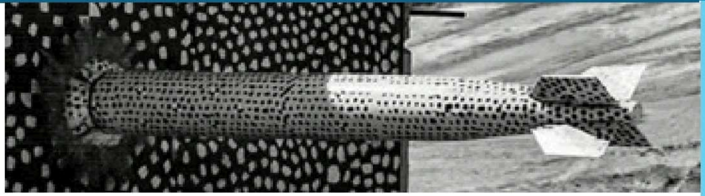

Ace Sorensen, Birk Jones, Ross Guttromson, Stephen J. Verzi
Sandia National Laboratories

Sandia Machine Learning and Deep Learning Workshop
August 6-7, 2019

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.

SAND2019-7862 PE

Abstract

The U.S. electric power grid is vital to daily function of all critical infrastructure sectors and directly impacts our daily standard of living. It currently faces a complex set of challenges as it continues to age with respect to technology modernization, environmental changes, and physical risks. Currently, for managing the US electric power grid, best practice utilizes extensive planning and is limited to maintaining narrowly-defined operational boundaries. Blackouts result from deviation outside the preplanned boundaries due to unforeseen interactions. System restoration, at that point, is based on intuition and experience of system operators. The identification of system vulnerabilities and remediations are determined by subject matter experts, but system complexity severely limits this effectiveness.

In the infrequent occurrence when grid operations depart from planned criteria, how do we move to a 'good' operating point? During near blackout conditions, grid operators may have an opportunity to restore the system to a safe condition if a real-time decision support tool is available. This project investigates the development of a real-time decision support tool for that purpose.

We explore coupling modeling and analysis methods from multiple domains to provide real-time decision support for mission critical infrastructure.

Our research leverages reinforcement learning employing deep neural networks (DNNs), as in AlphaZero (Silver et al., 2018), to identify "best" (or approximately optimal) resilience strategies for operation of a grid model. Current results have demonstrated the potential for learning in a model of grid state space navigation, but continued work is needed.

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Project Objective and Purpose

During near blackout conditions, grid operators may have an opportunity to restore the system to a safe condition if a real-time decision support tool is available.

This project investigates the development of a real-time decision support tool for that purpose.

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Existing Planning and Operations

1. In the infrequent occurrence when grid operations depart from planned criteria, how do we move to a 'good' operating point?
2. Where are we? Where do we want to go? What path do we take?

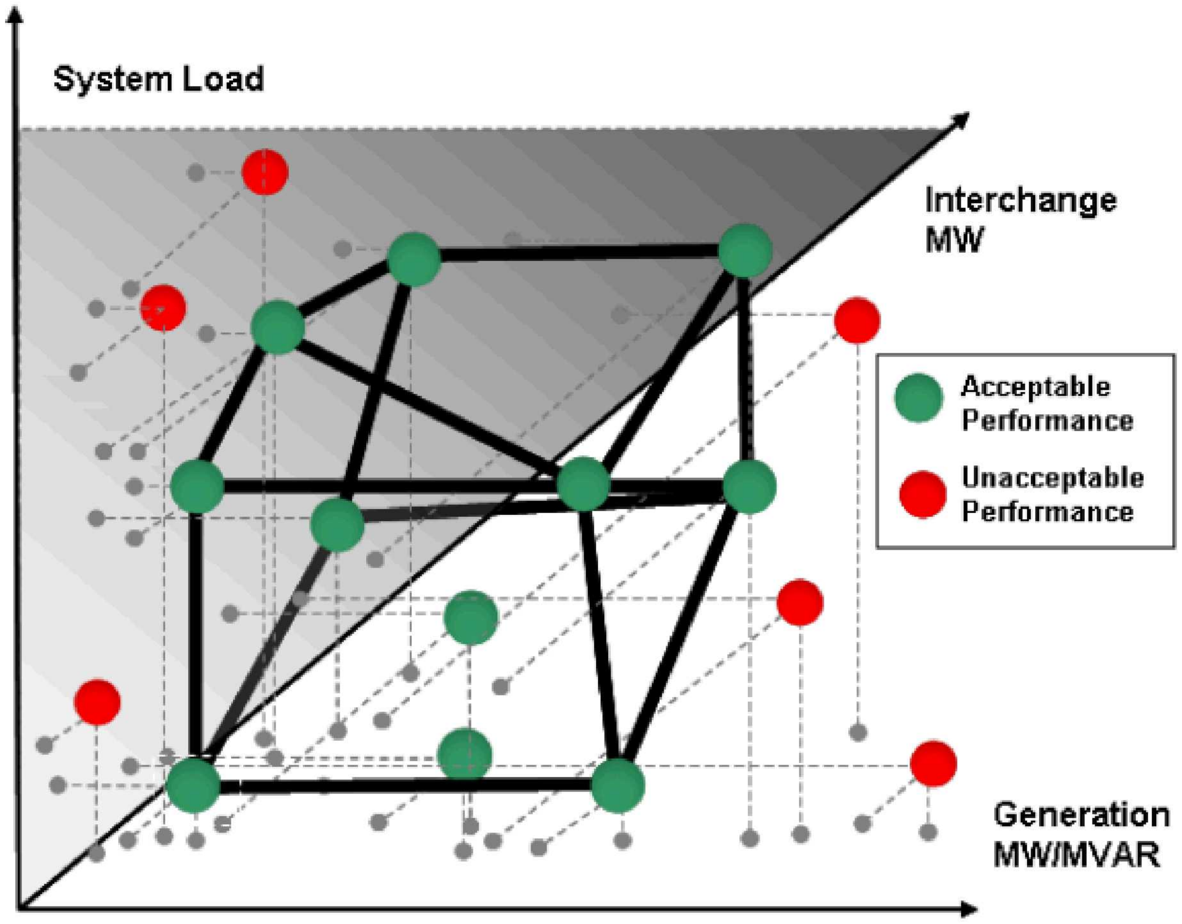

Figure 9 - "Scatter" plot of planning scenarios.

Metaphor For Stability Margin

Require "Stability" Margins of Interest

- Voltage Stability Margin
- Transient Stability Margin
- Non-Linear/Eigen-analysis Stability Margin
- System Voltage Margins
- Power Line Transfer Margins
- System Droop Margin

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

General Method

The solution method uses Deep Neural Networks combined with Monte Carlo Decision Trees to represent the sequence of control actions and dispatches needed for the grid to increase its stability margins.

Our approach offers the potential of a speedy solution to this problem, with a low risk of non-convergence. The solution will not be proven optimal, although it will be demonstrated to be feasible and 'good' during off-line testing.

Based on recent work conducted at Sandia

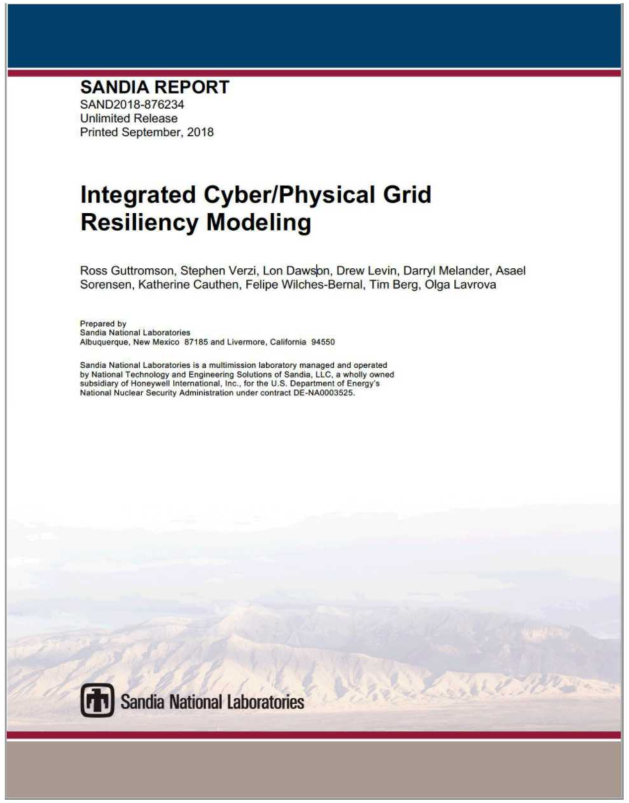

SANDIA REPORT
SAND2018-876234
Unlimited Release
Printed September, 2018

Integrated Cyber/Physical Grid Resiliency Modeling

Ross Guttromson, Stephen Verzi, Lon Dawson, Drew Levin, Darryl Melander, Asael Sorensen, Katherine Cauthen, Felipe Wilches-Bernal, Tim Berg, Olga Lavrova

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

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.

Sandia National Laboratories

Defining the State Space

Prob

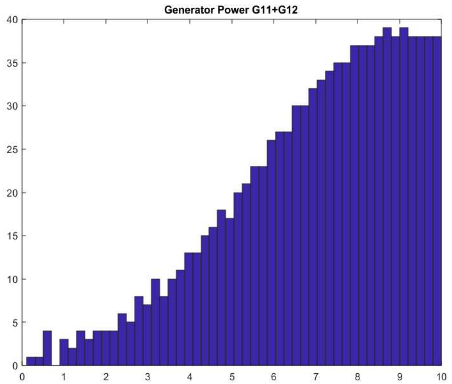

Pgen, pu

Prob

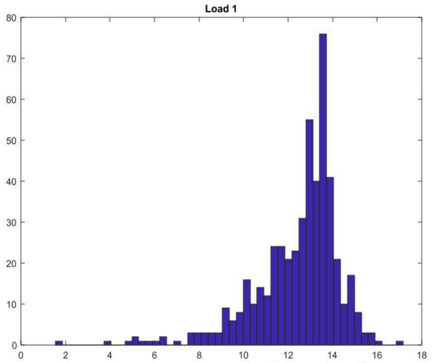

Load, pu

Prob

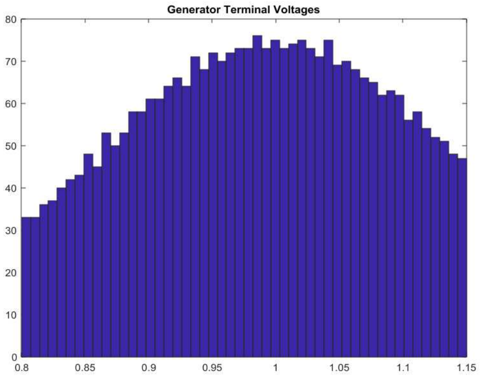

Vt, pu

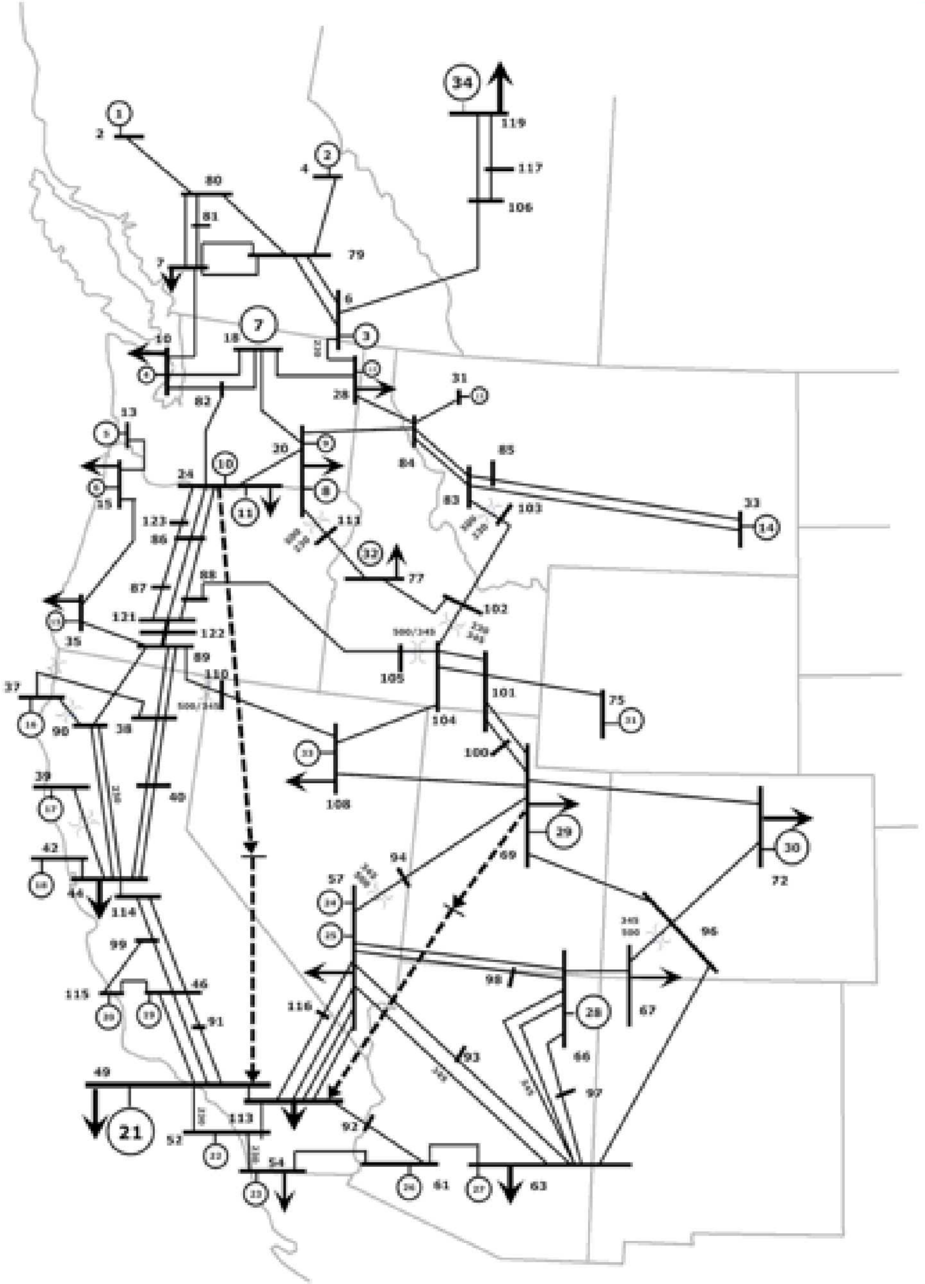

Stability Margins

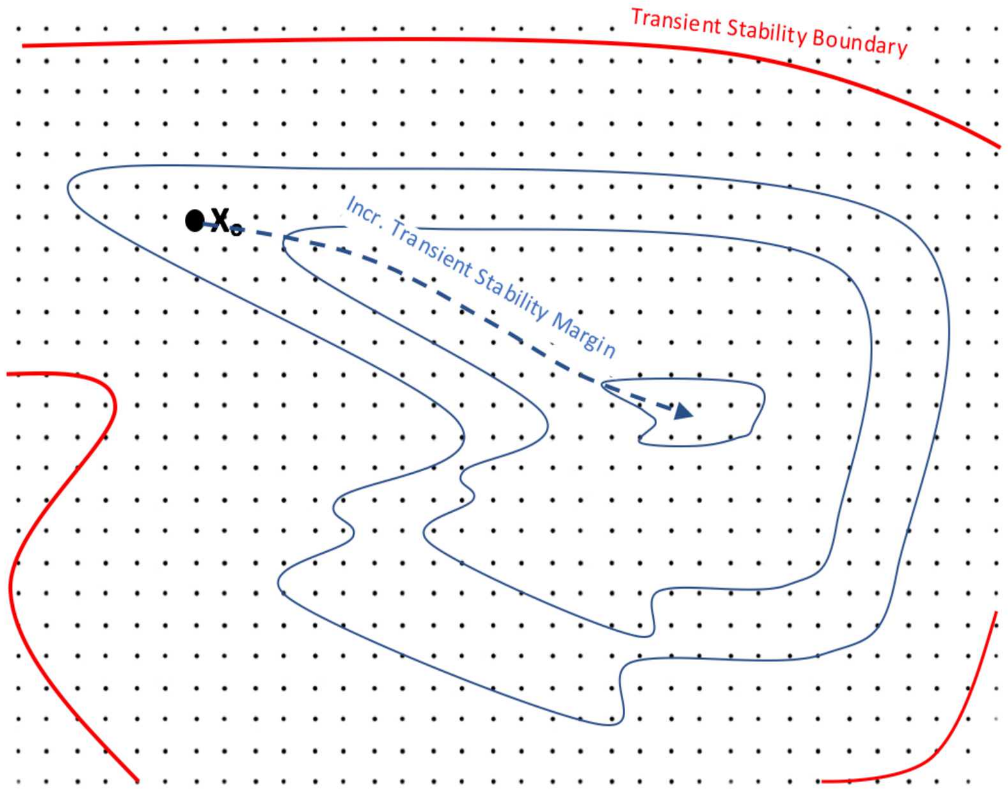

Transient Stability Level Curves

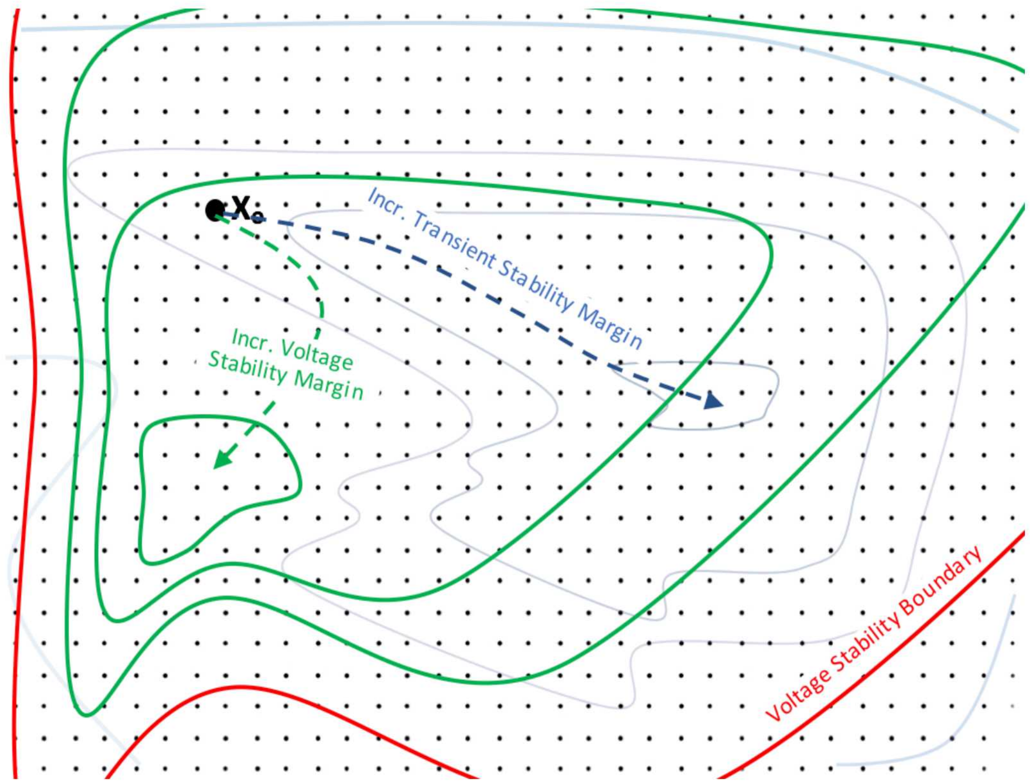

Transient Stability and Voltage Stability Level Curves

The AlphaGrid Game

- Game objective
  - Increase score by moving from the current state to a state with high stability margins
    - In the fewest moves possible
    - While maintaining high stability margins during all state changes
- Game setup
  - Game board is the grid state space and associated stability margin penalties
  - Initial state is a marginally stable grid condition
- Scoring includes
  - An aggregate of stability margin penalties along the journey of grid states
  - Penalties for more transitions
  - Improves for a transition to a more stable state
  - Degrades for a transition to a less stable state
- Rules for discrete state transitions
  - Each state transition is bounded to nearby neighbors
  - Journey to the final state is performed using a sequence of state transitions
  - States are not allowed to be re-visited
  - Cannot transition to unstable state
  - State transitions are
    - Selected from the combined DNN and MCTS during game play
    - Learned during training in the DNN
- End of the game is reached when
  - The maximum number of transitions is exhausted
  - No possible transition to stable states exists
  - Sufficient stability margins have been reached

AlphaGrid Deep Neural Network Block Diagram

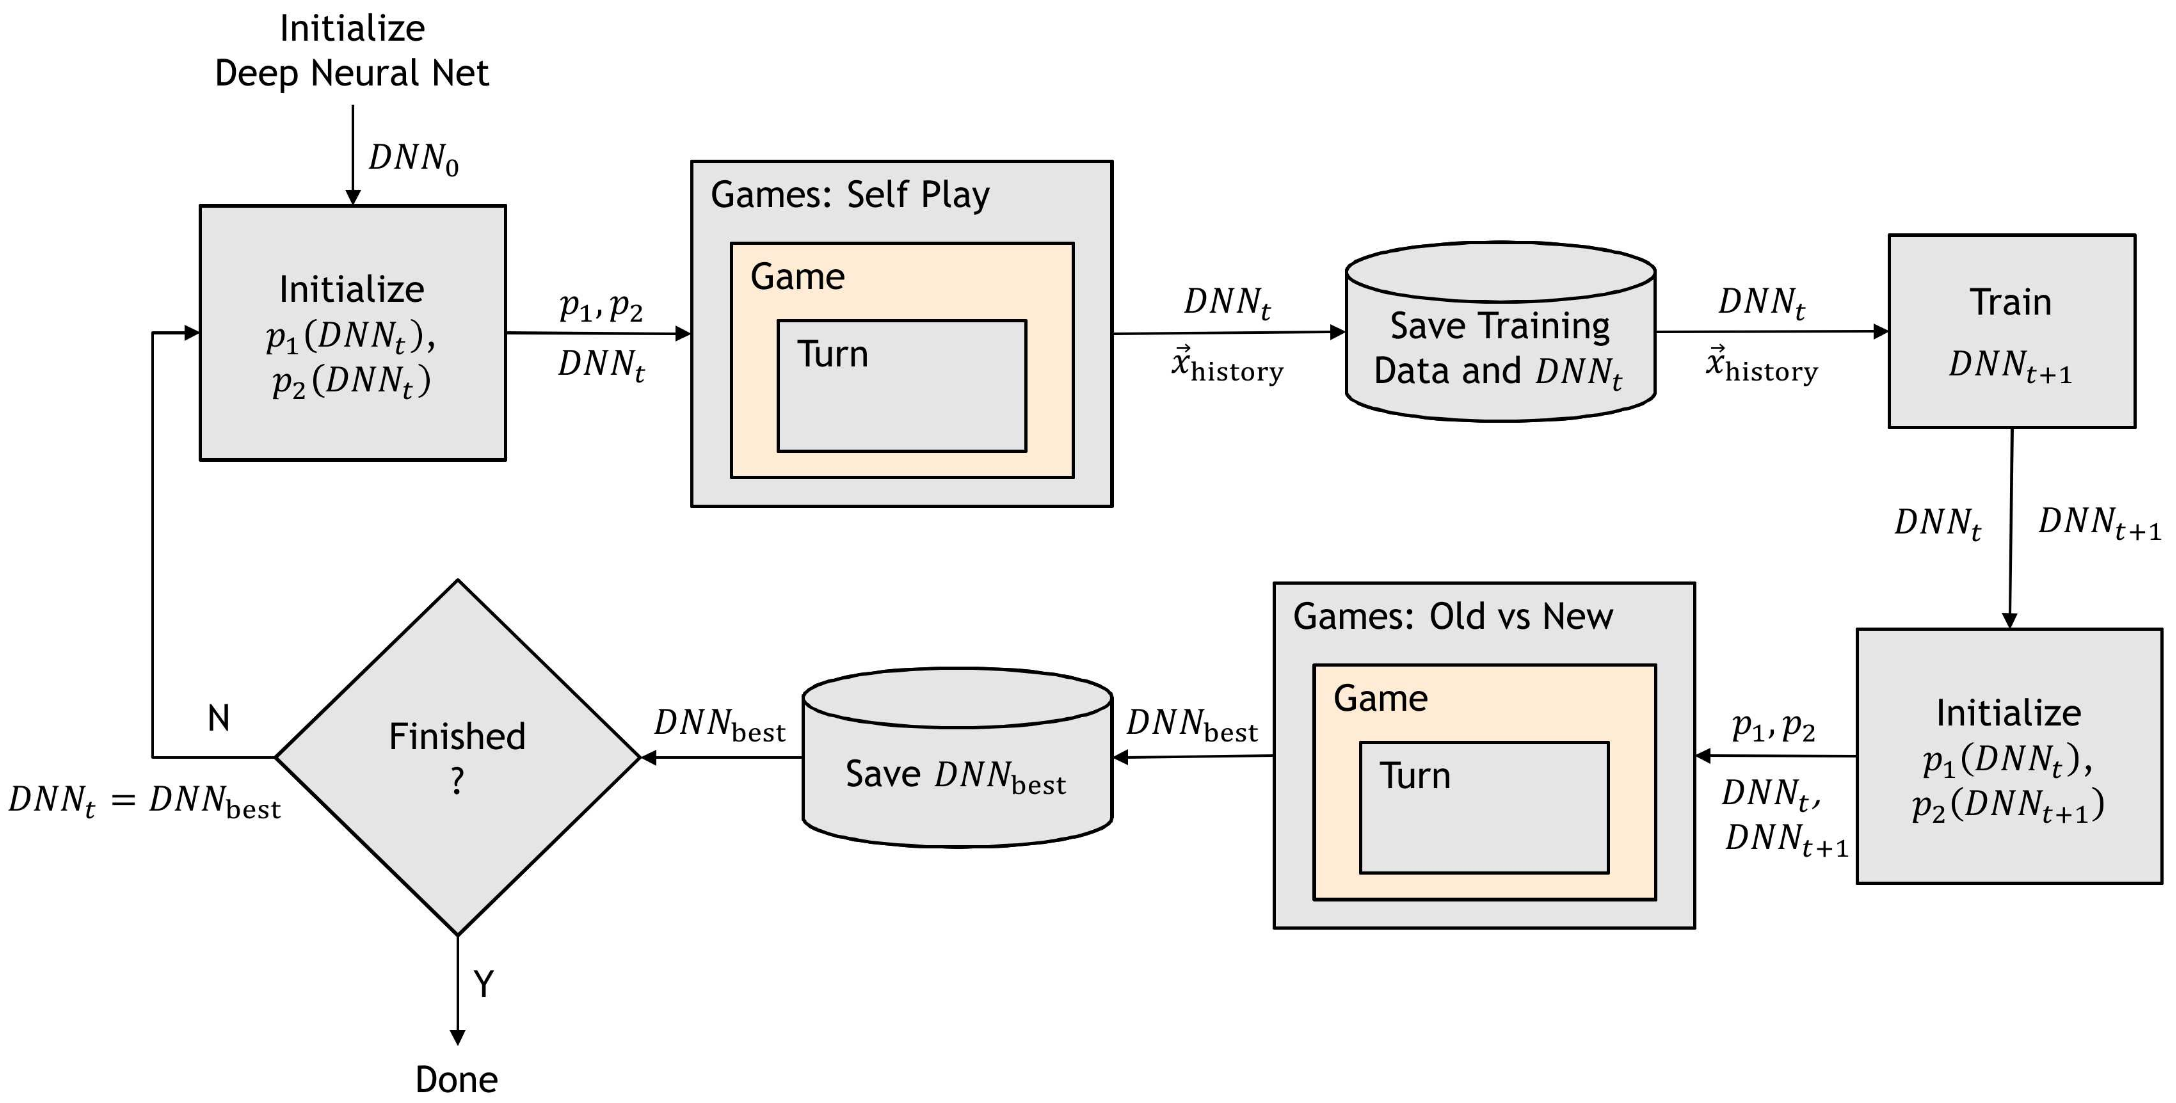

AlphaGrid Deep Neural Network Block Diagram

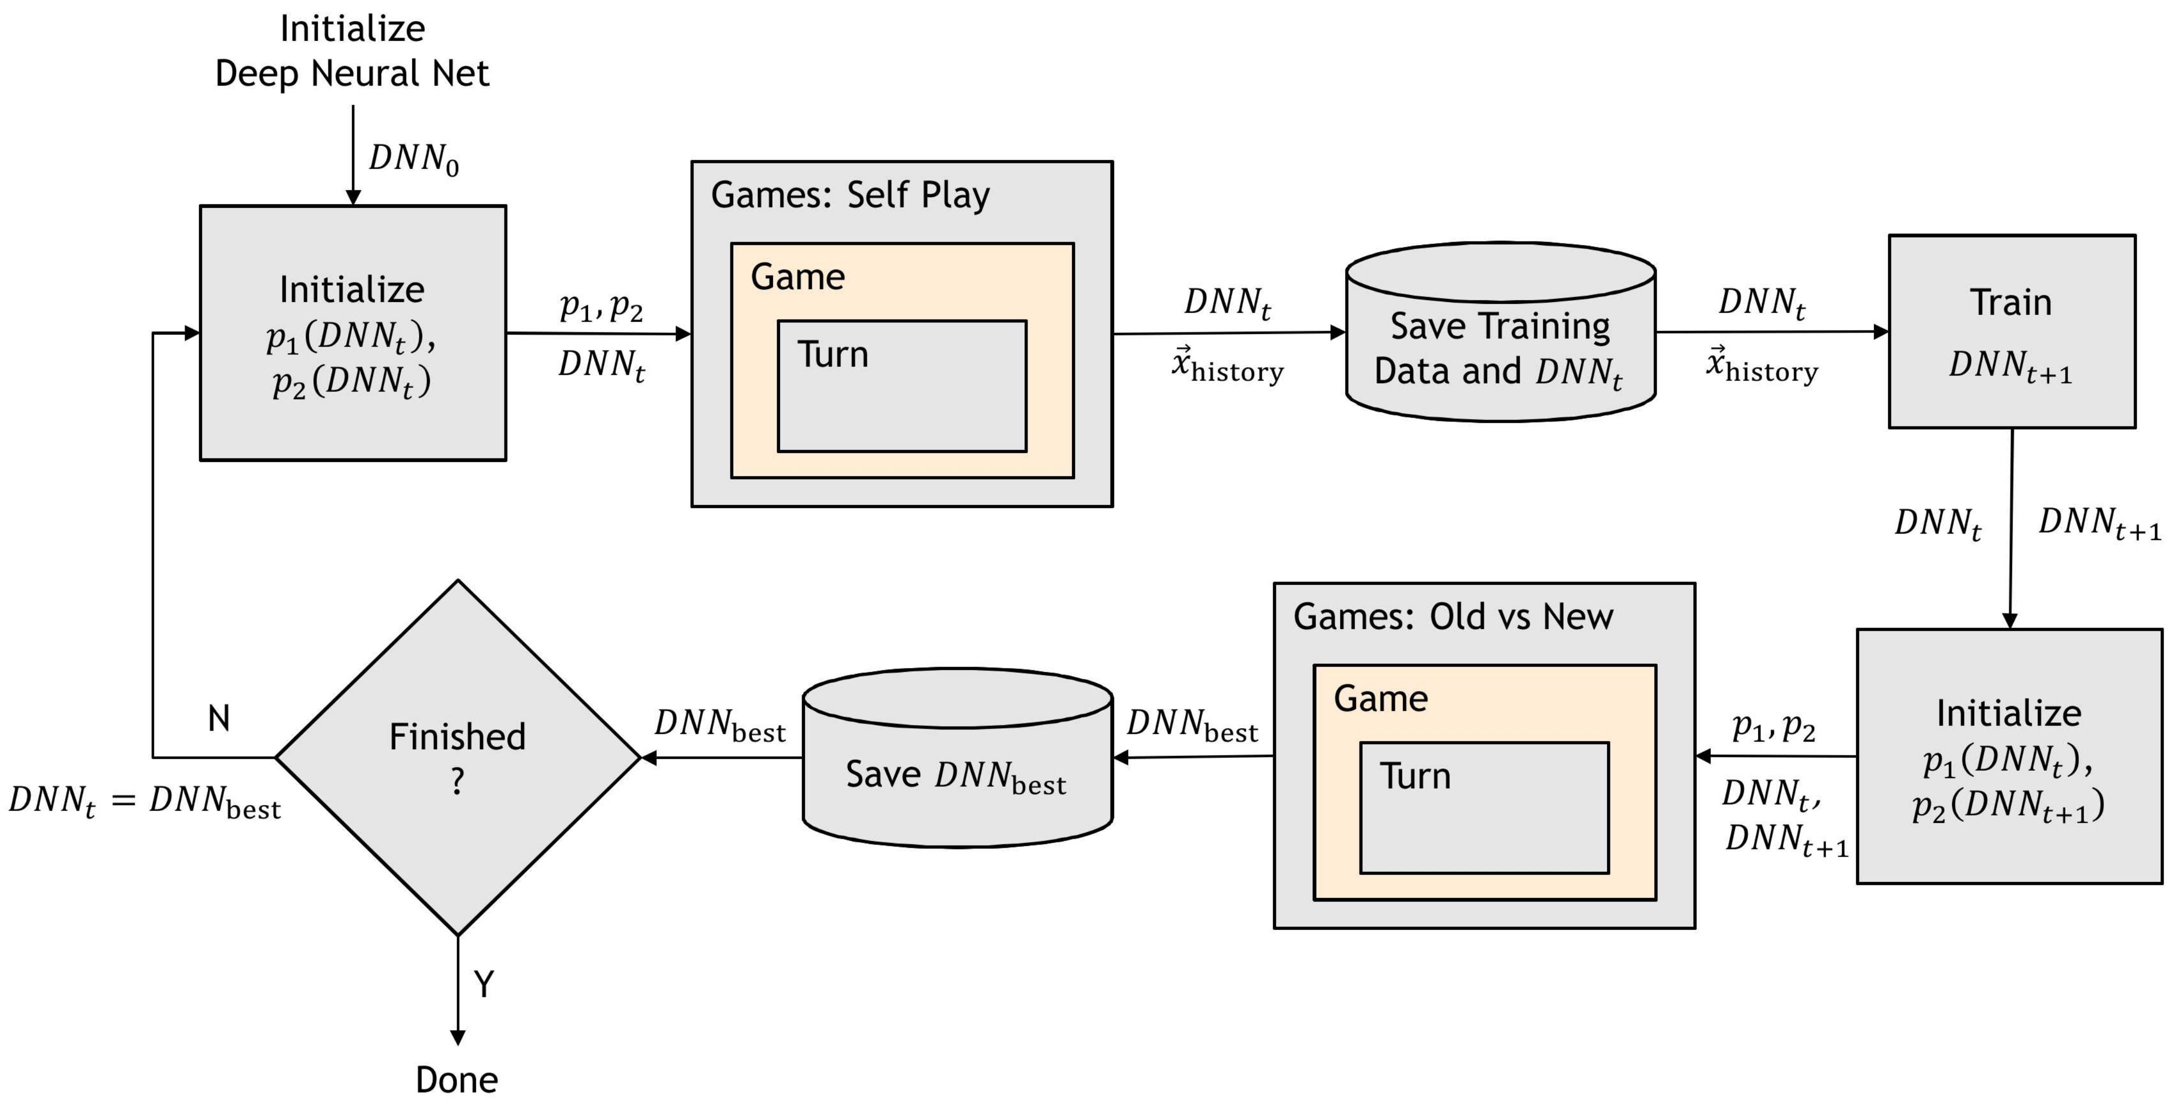

Game Play Block Diagram

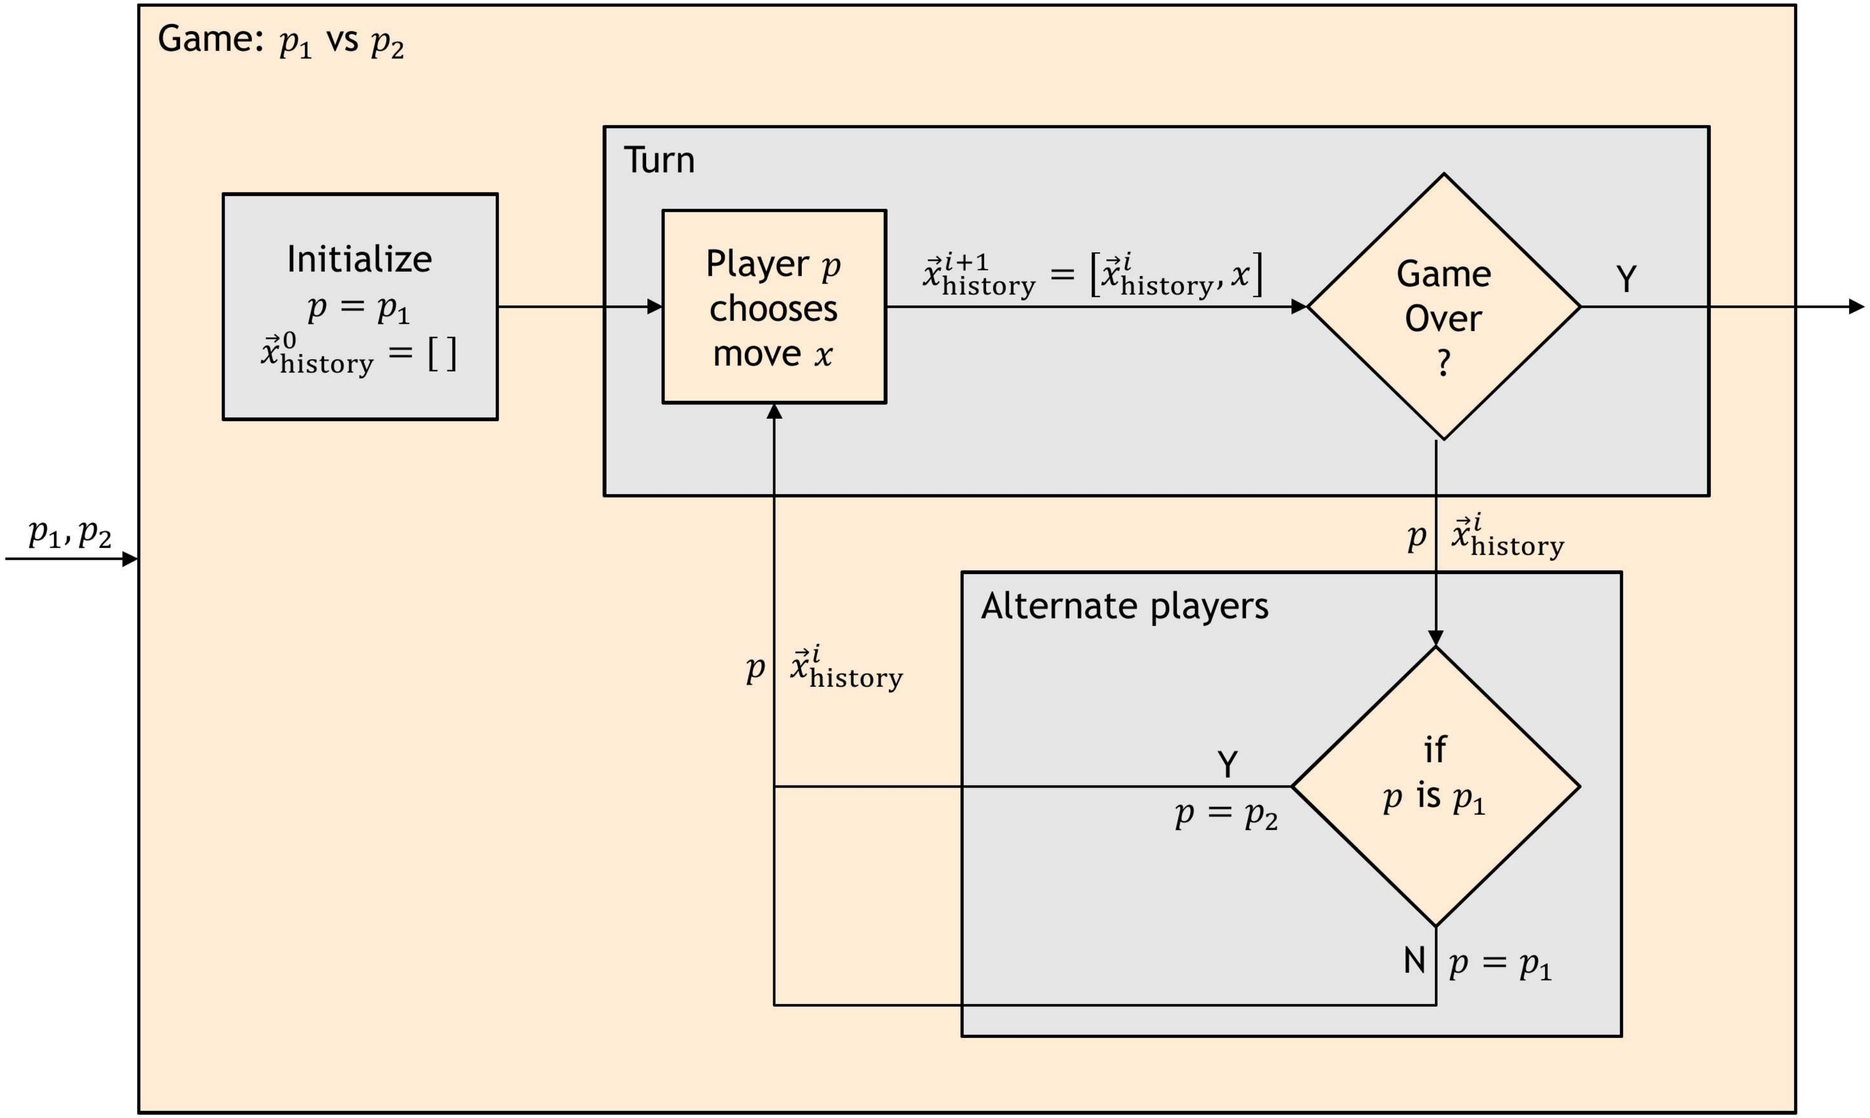

Monte Carlo Tree Search (MCTS) Block Diagram (Game Play Turn)

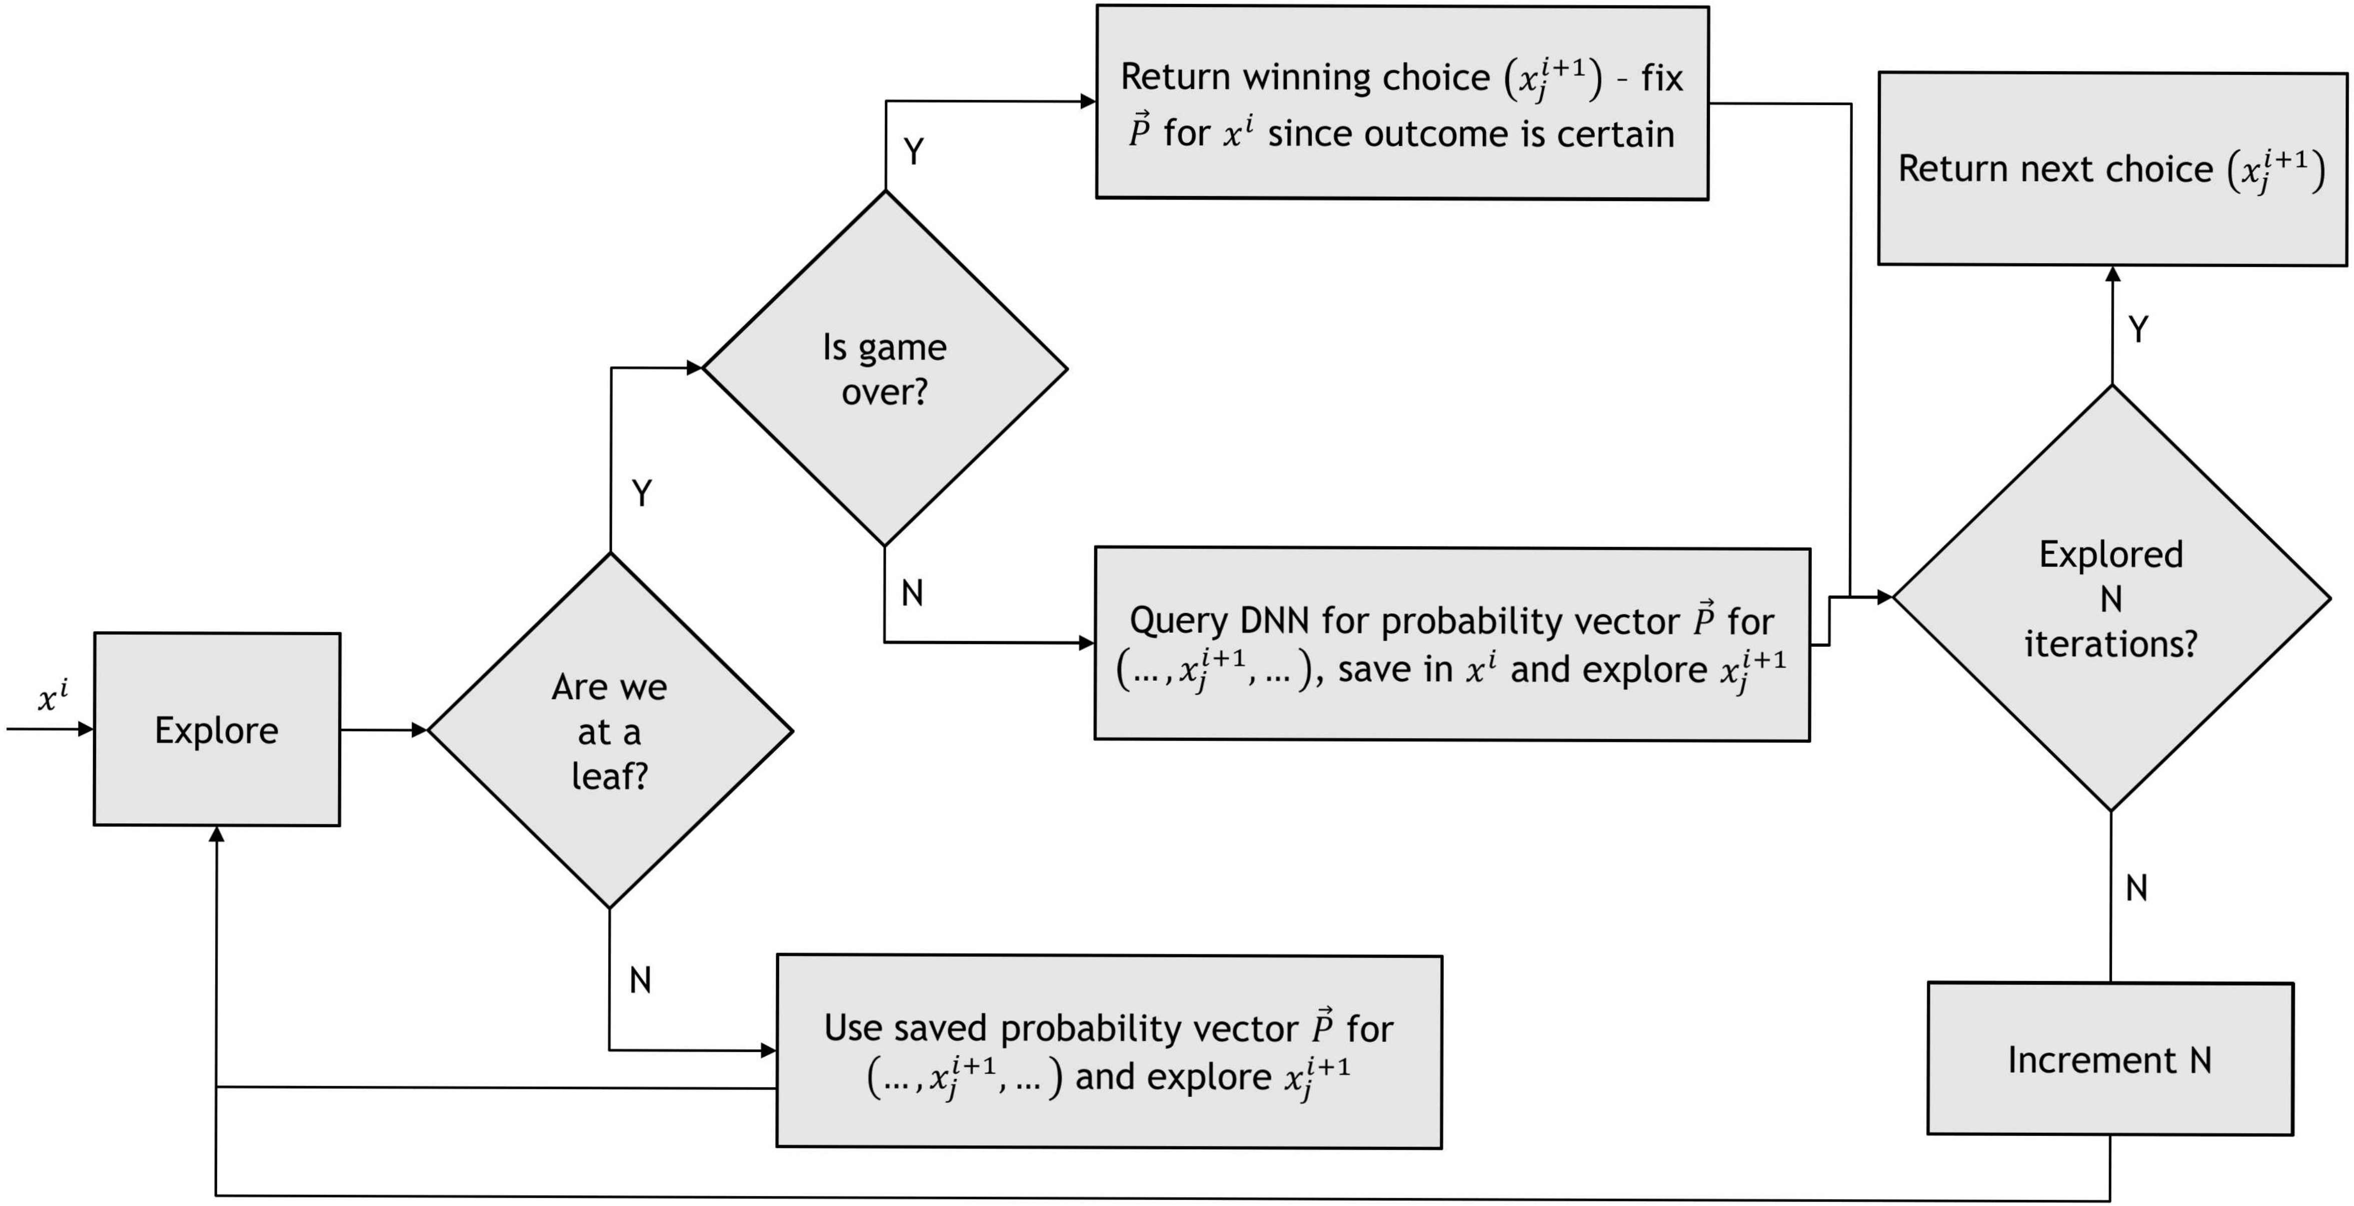

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

AlphaGrid Learning Results

Grid model operator comparison

- Trained vs Random – average 89.6% of the time Trained DNN out-performs Random player (in last half of training iterations)

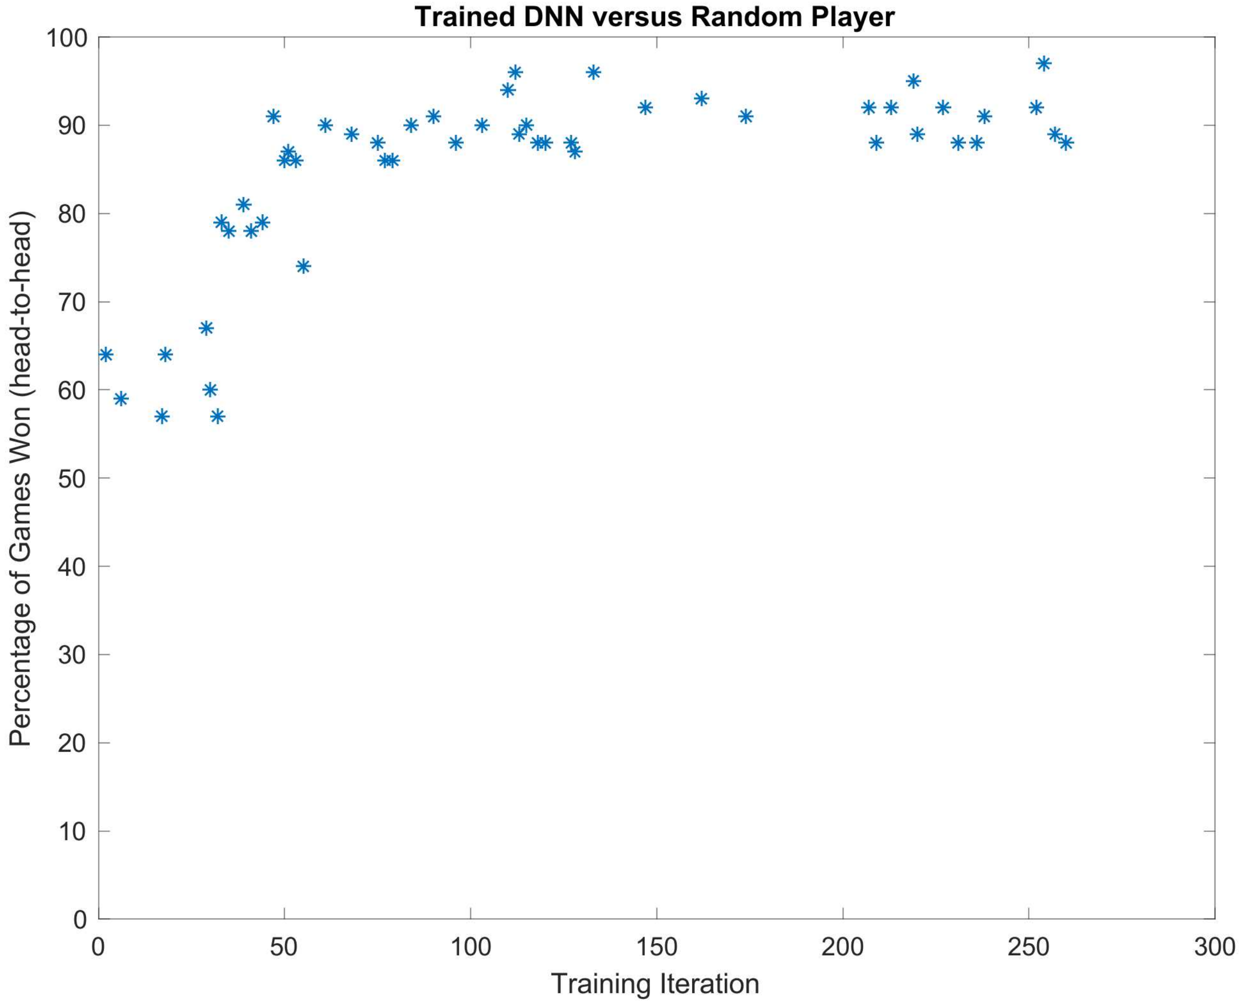

Random player refers to the choice of next state from current, where a random choice (without repeat) reachable from the current state is chosen

AlphaGrid Learning Results

Grid model operator comparison

- Trained vs Greedy – average 26.6% of the time Trained DNN out-performs Greedy player (in last half of training iterations)

Greedy player refers to the choice of next state from current, where the most stable next state reachable from the current state is always chosen

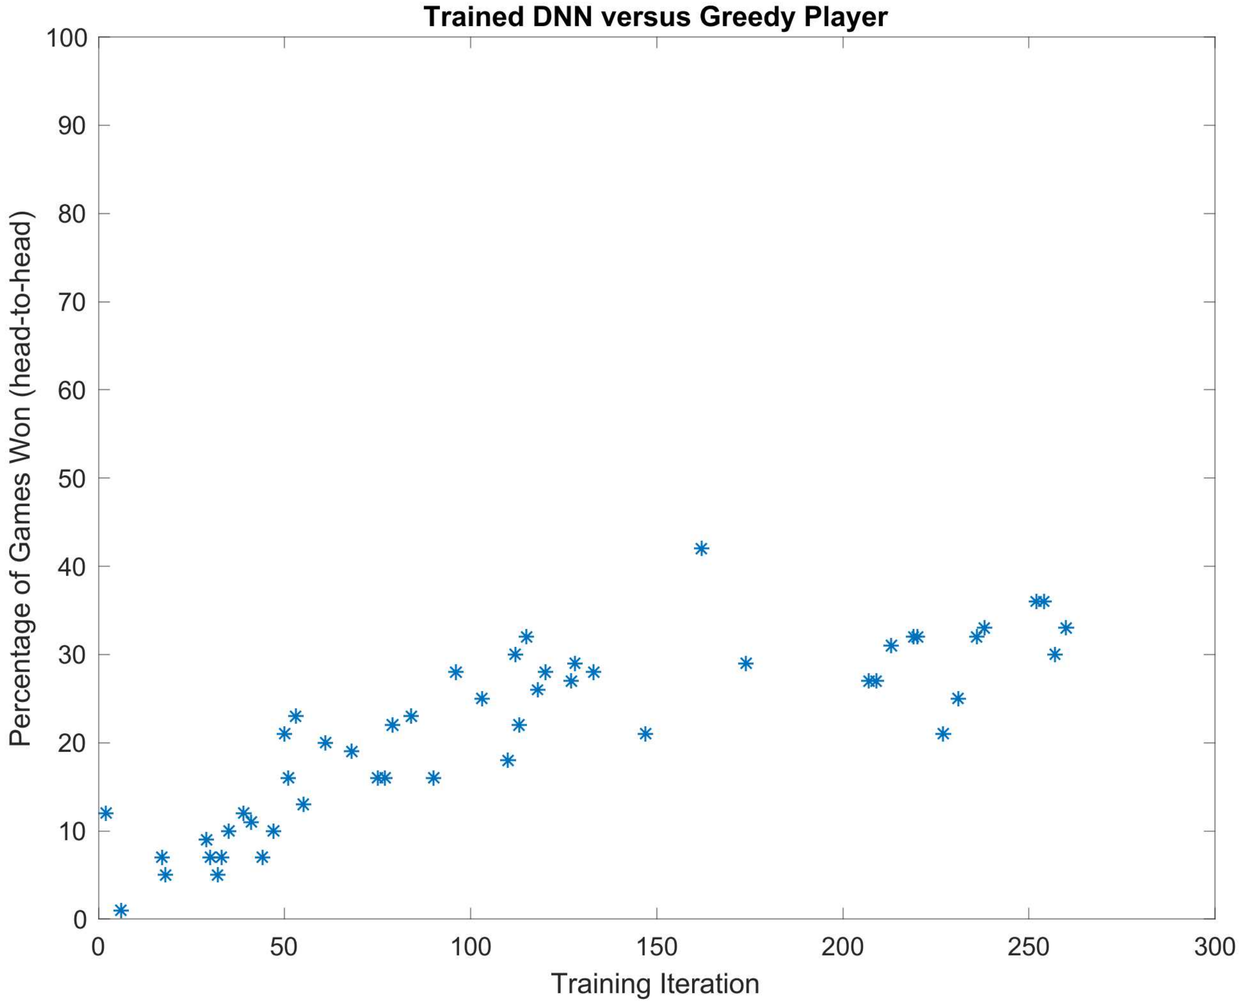

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Conclusions

- Demonstrated feasibility of approach
- Verified learning shows improvement in comparison to
  - Random state walk
  - Greedy state transition
- Just getting started

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Key Research Challenges

- Scalability of State Space Dimensions- Map reduction without losing fidelity
- DNN partitioning – Provides metadata, allowing insights into each solution (geography, type of constraint, etc)
- Modify solution from discrete space to continuous space
- Use of Transfer Learning- Allows a DNN to be trained for a specific system without starting from a blank DNN
- Management of cyclical state transitions
- Evaluation of tool across multiple scenarios- checking solution accuracy against many constraints

Project Goals for 2019

- A demo will be conducted, and a paper will be submitted for publication
- Construction and validation of a state space map at approximate control dimensionality of $\mathbb{R}^{20}$
- Demonstrate operation of the ML decision process using the reduced WECC system and perform solution verification using dynamic power system analysis.
- Start the development of state space compression methods.

Outline

- Project Purpose
- Significance & Impact
- Technical Approach
- Technical Accomplishments
- Conclusions/Summary
- Future Effort
- Acknowledgments and Contacts

Our Team – Sandia

- Ross Guttromson (co-PI)
- Stephen J. Verzi (co-PI)
- Christian "Birk" Jones (grid modeling)
- Asael "Ace" Sorensen (deep reinforcement machine learning)
- Raymond "Ray" Byrne (PM)
- Charles "Charlie" Hanley (Senior Manager)

Thank You

Questions?

Backups

Approach: Machine Learning

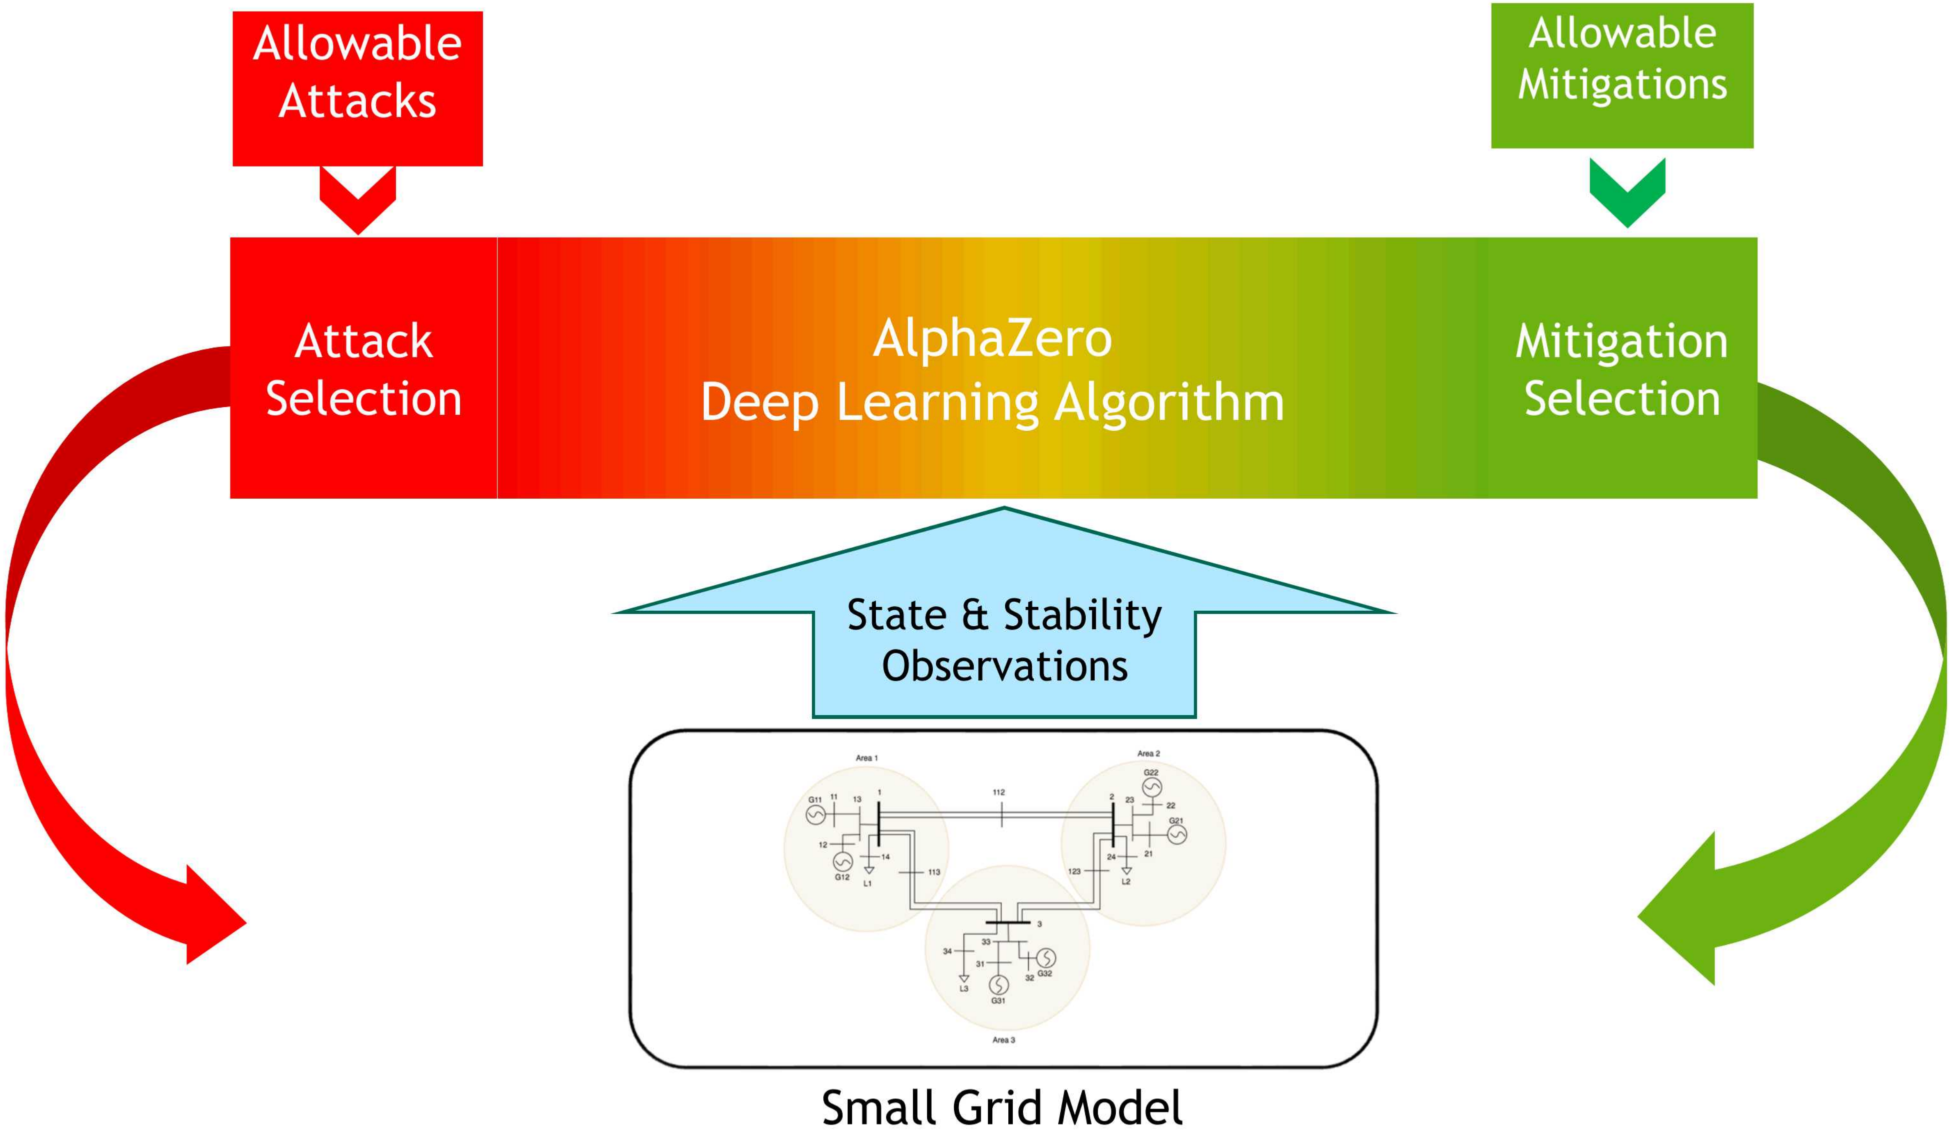

Performance

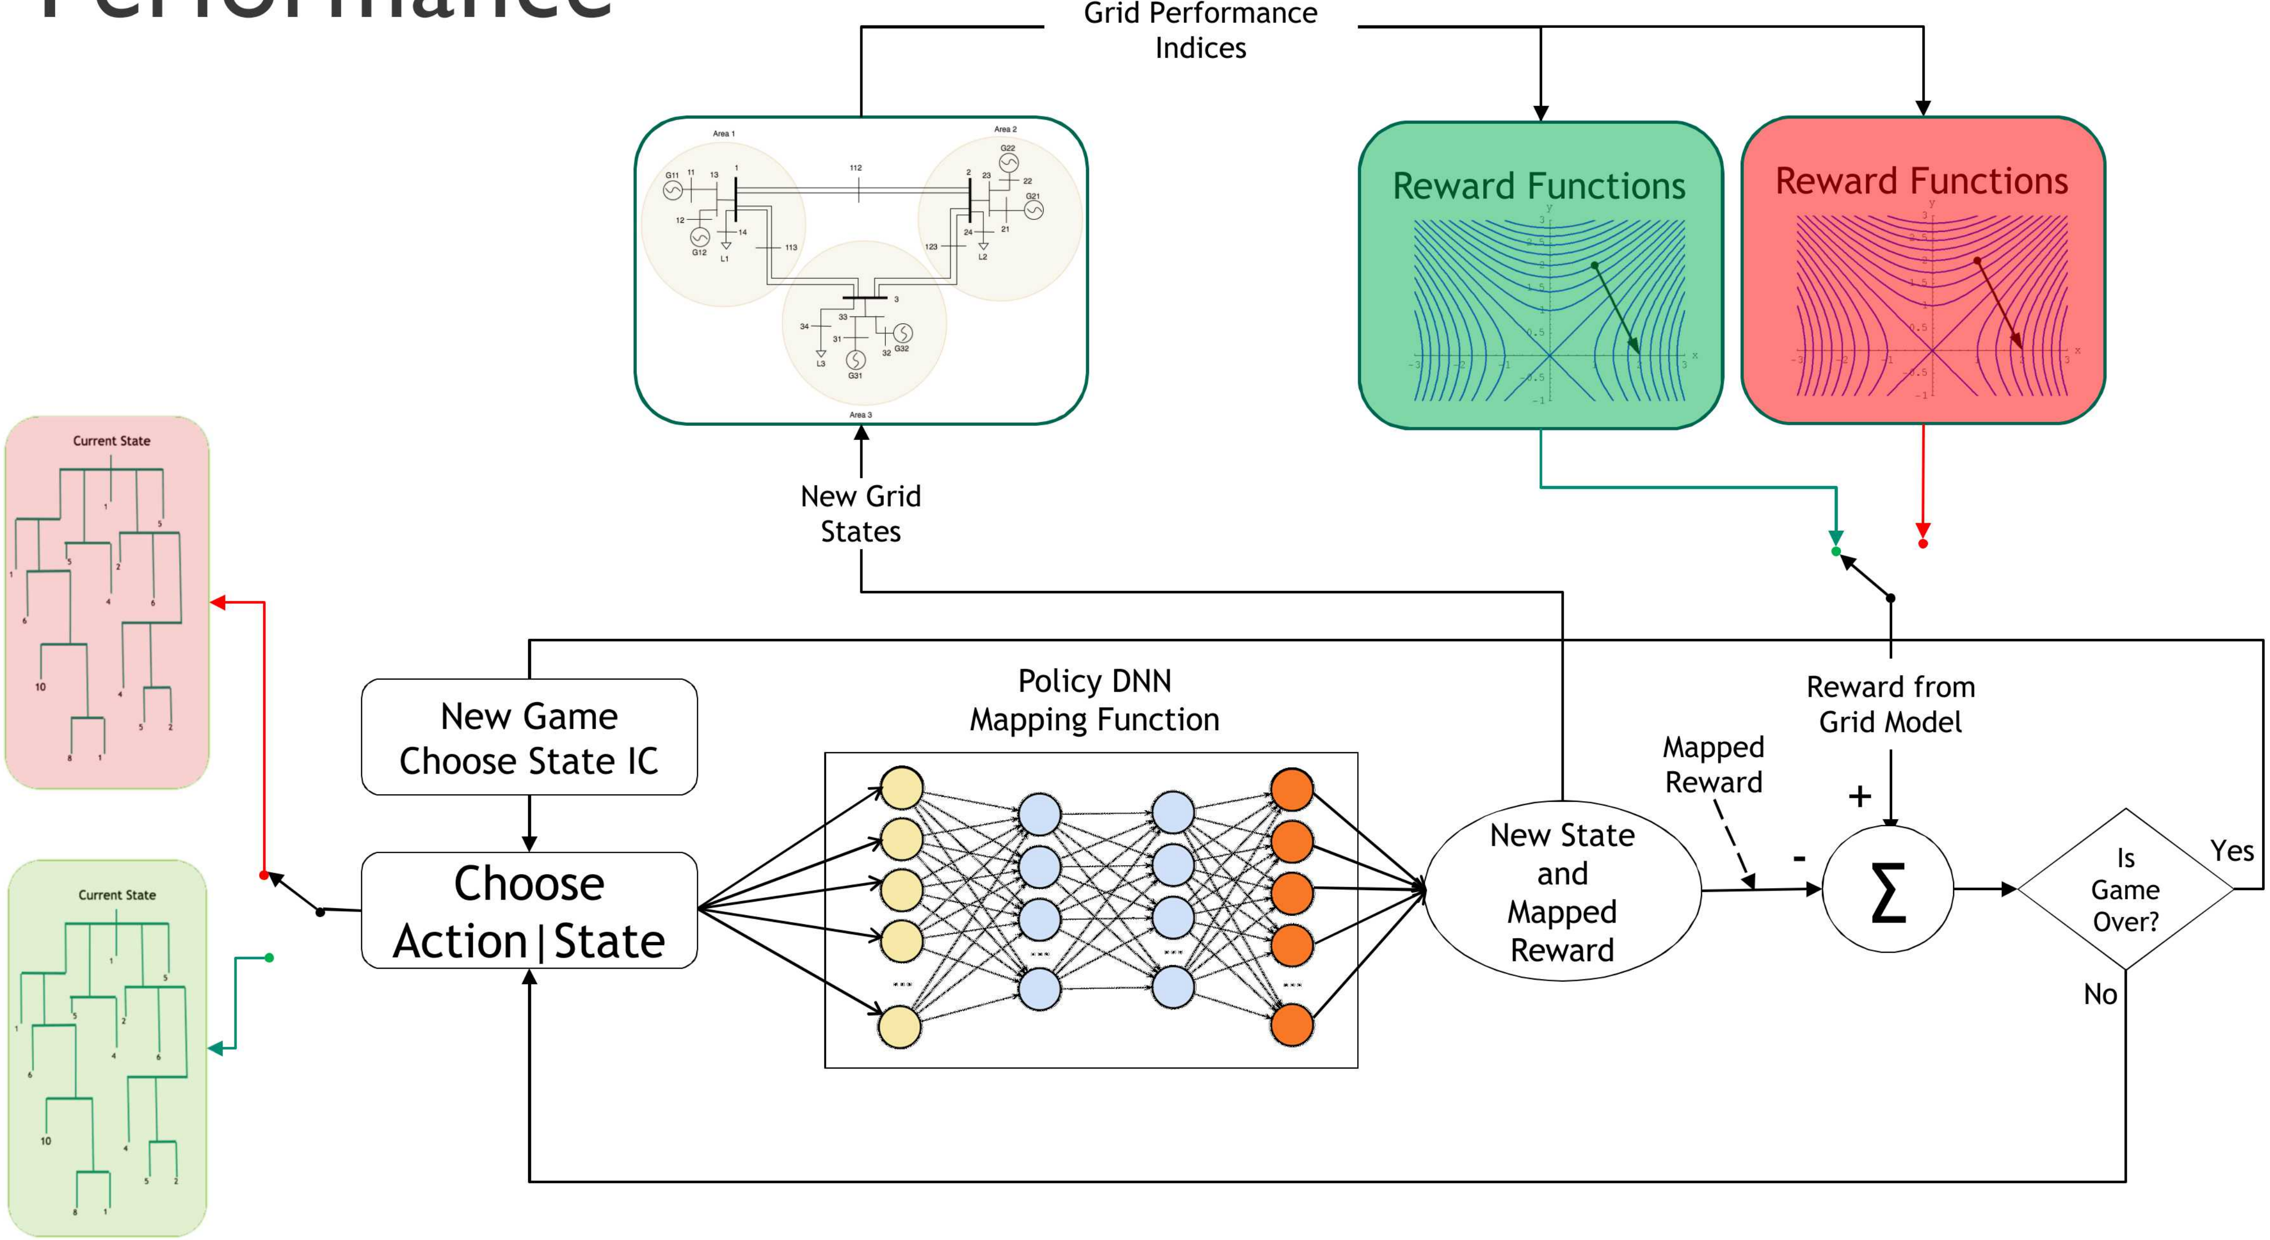

AlphaGrid Deep Neural Network Block Diagram (one-player)

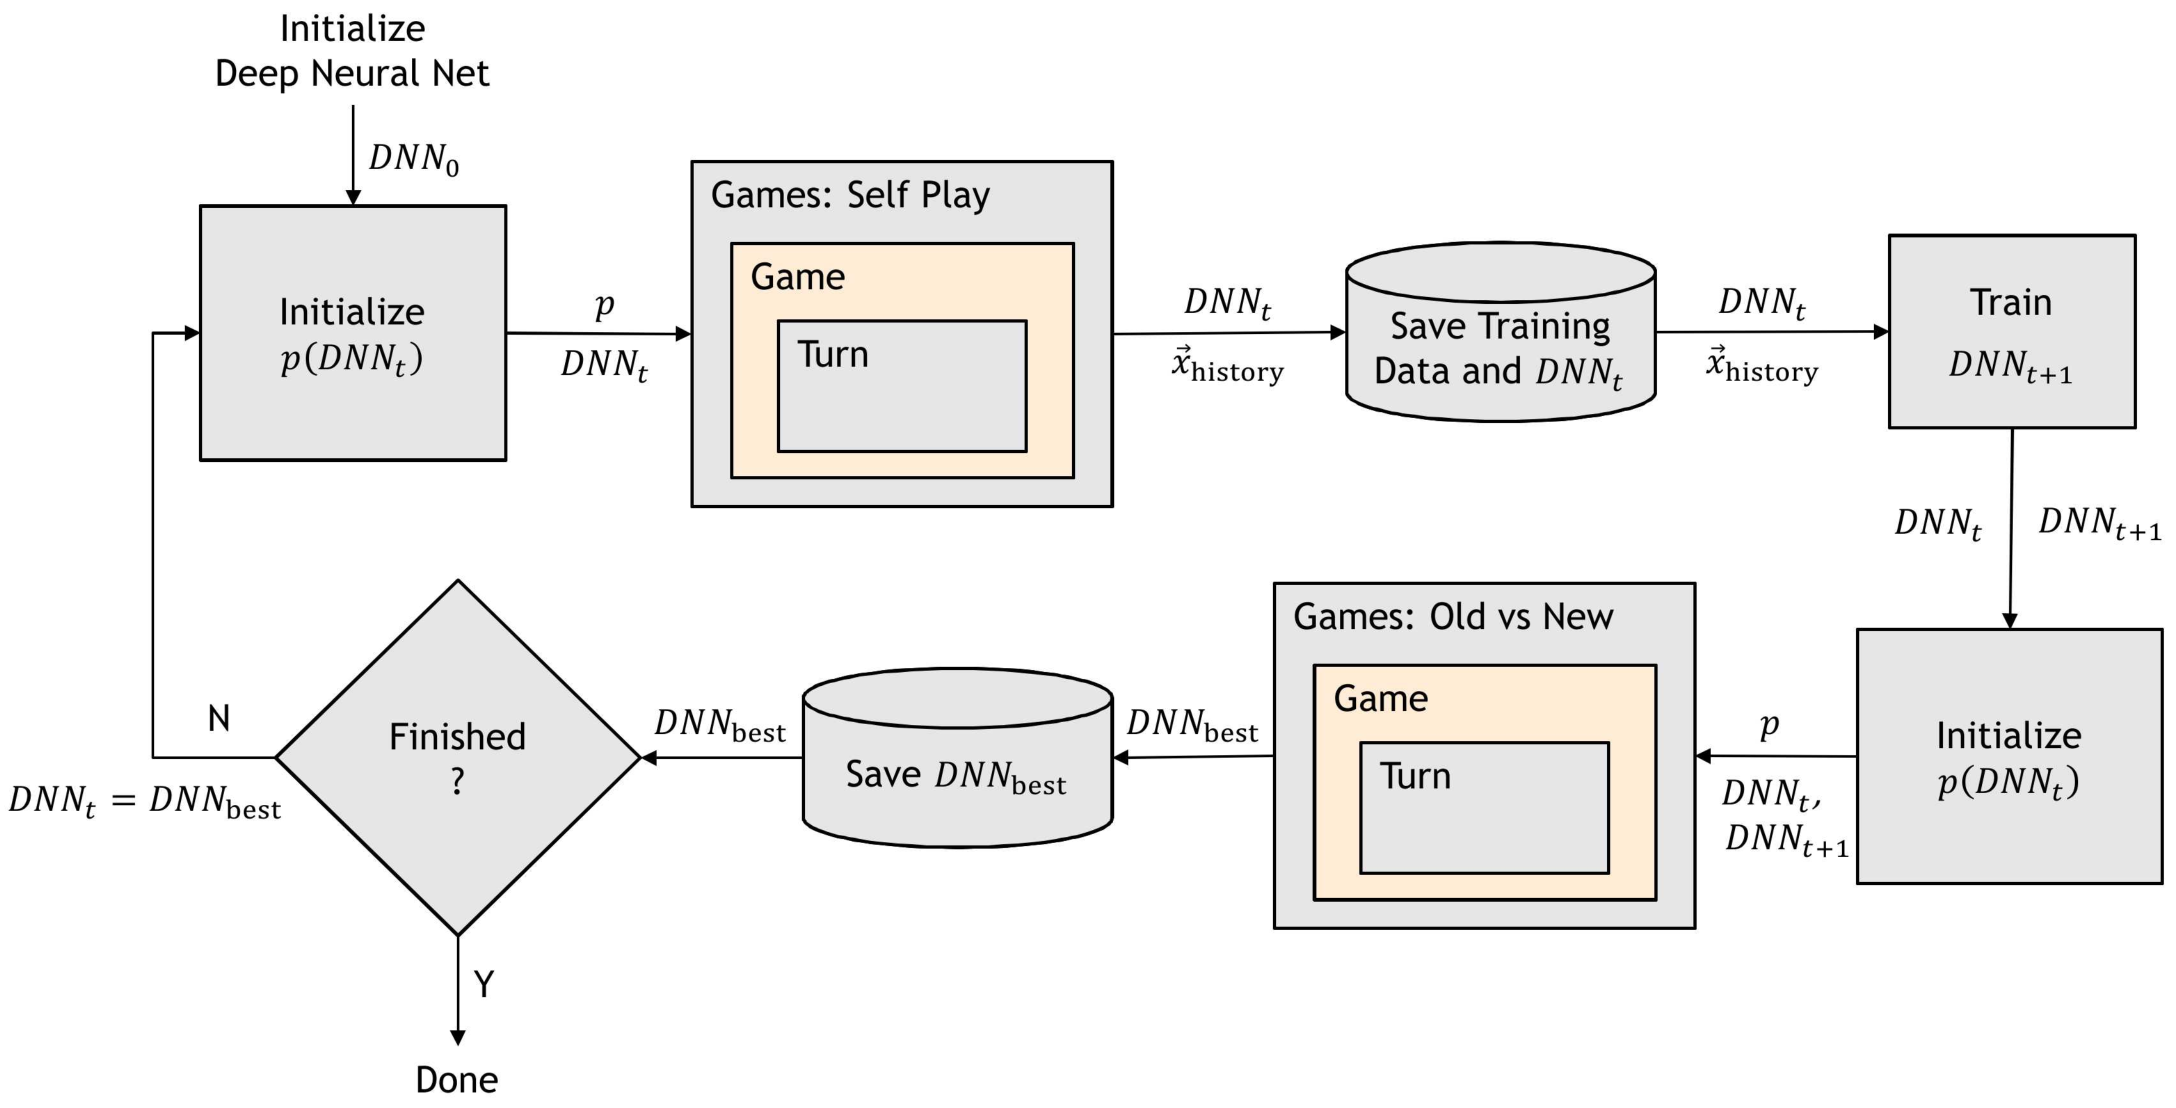

Structure (of Grid State Data) infers validity of Machine Learning Approach

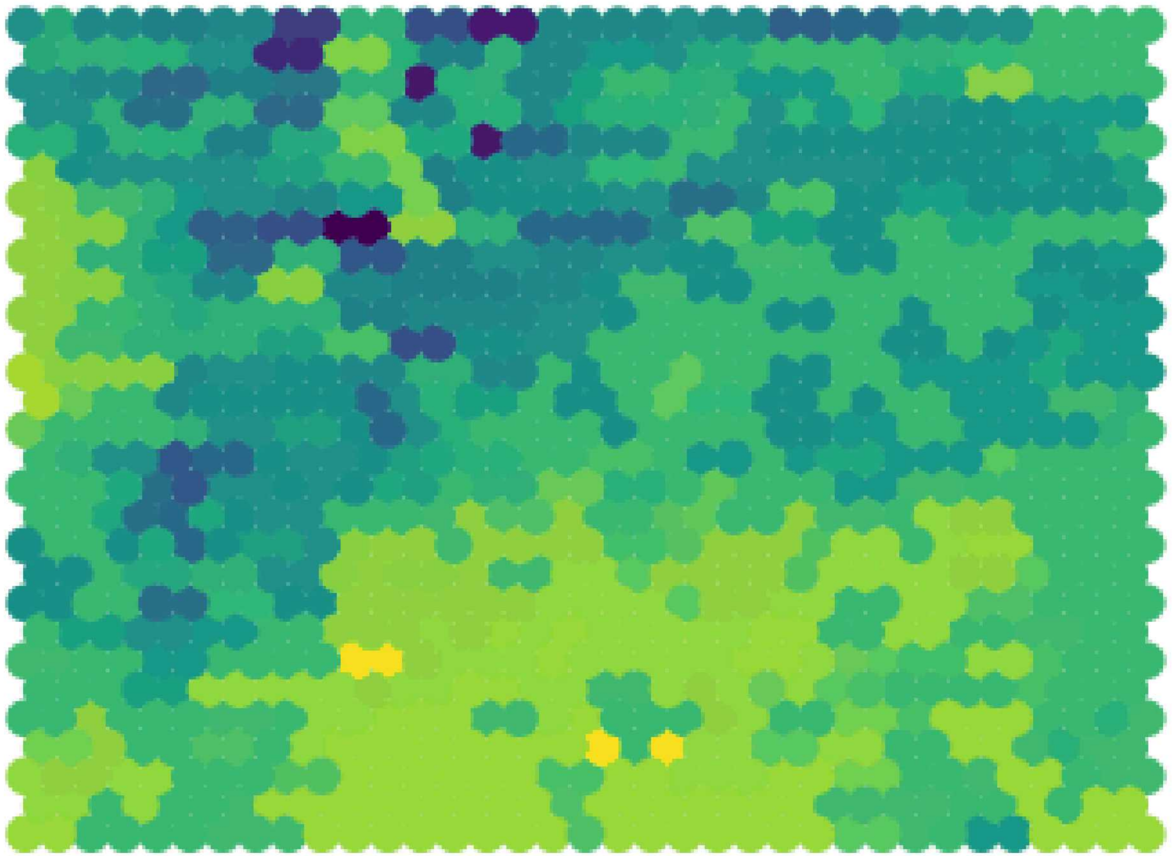

15-dimensional grid state data of 1,001 precomputed points flattened onto a 2D hex grid for visualization. Light yellow represents high stability scores, dark blue represents low stability. The plot shows spatial correlation and bounded stability regions, validating a machine learning approach!

AlphaGrid – playing the game (attacker)

On its own an attacker would play a trajectory from good (yellow) to bad (blue)

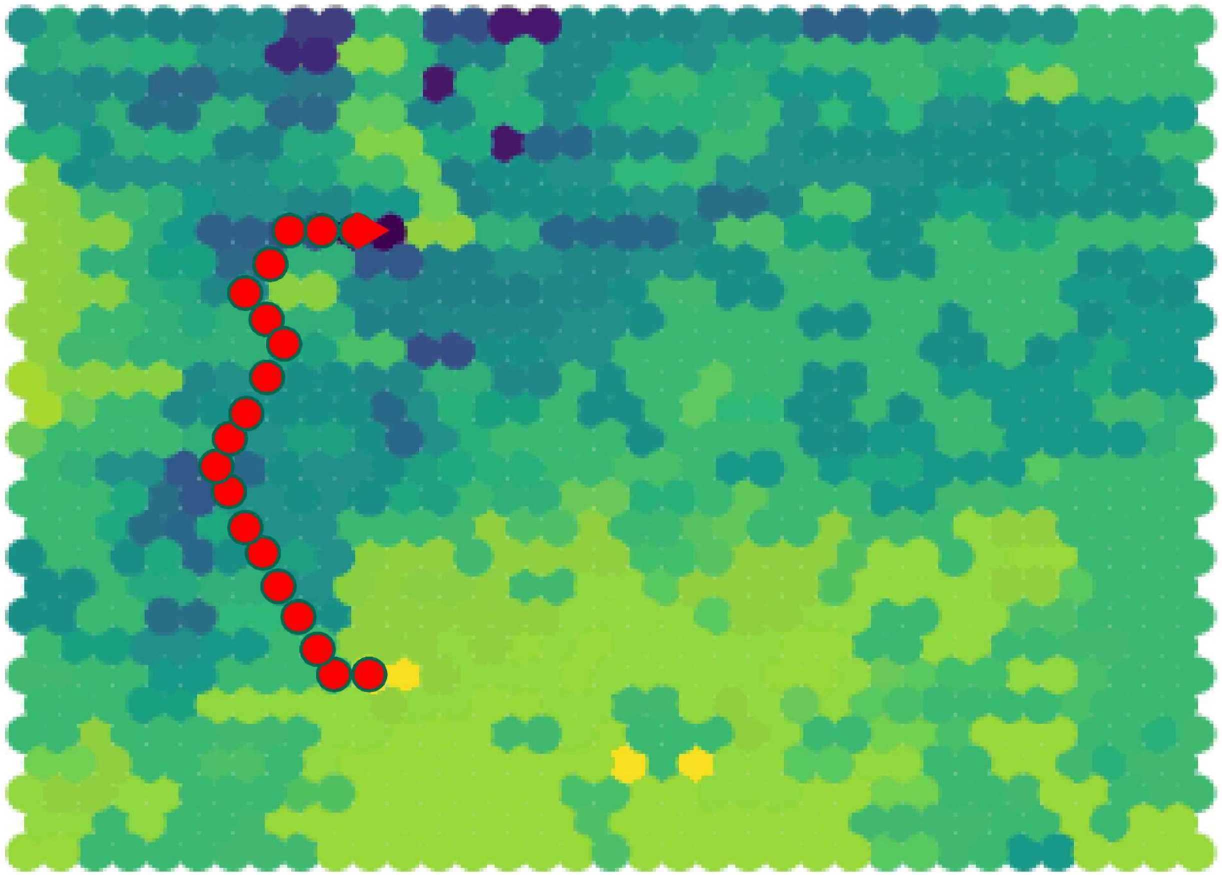

AlphaGrid – playing the game (defender)

On its own a defender would not necessarily play a trajectory from bad (blue) to the nearest good (light yellow) but rather would choose good within a group of other good (yellow within a sea of yellow)