Deploying High Performance Computing Applications of AI/ML in the Energy Sector

Nefeli Moridis, Ph.D.
NVIDIA

Vidyasagar, Ph.D.
Beyond Limits

AI and Accelerated Computing

Robotics
Manufacturing, construction, navigation

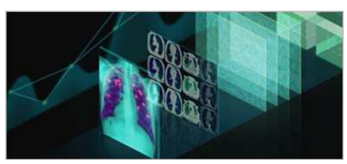

Healthcare
Cancer detection, drug discovery, genomics

Autonomous Vehicles
Pedestrian & traffic sign detection, lane tracking

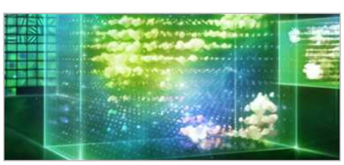

Internet Services
Image classification, speech recognition, NLP

Media & Entertainment
Digital content creation

Intelligent Video Analytics
AI cities, urban safety

At the Intersection of Graphics, HPC, and AI

Webinars

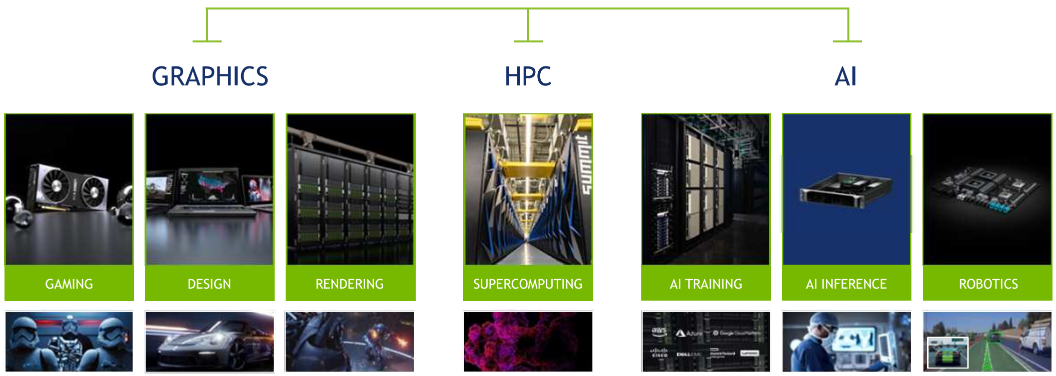

SPE International

The Realm of Artificial Intelligence

Webinars

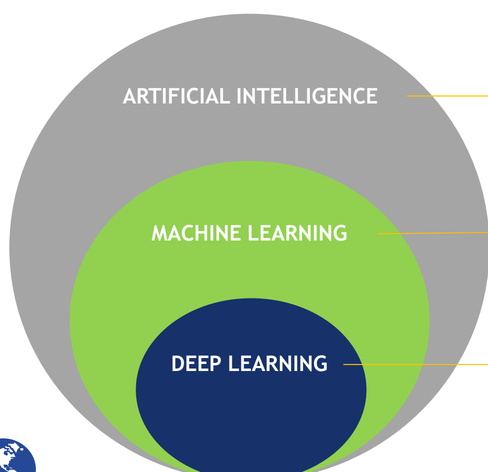

Solves a task requiring human intelligence
Any technique that allows computers to mimic human behavior

Solves a specific AI problem by training a machine to learn from data
Uses well-researched statistical methods and algorithms to train predictive models

Solves an ML problem by using multi-layer neural networks
Particularly powerful when dealing with complex inputs like images, text & speech

SPE International

Myths and Truths

Webinars

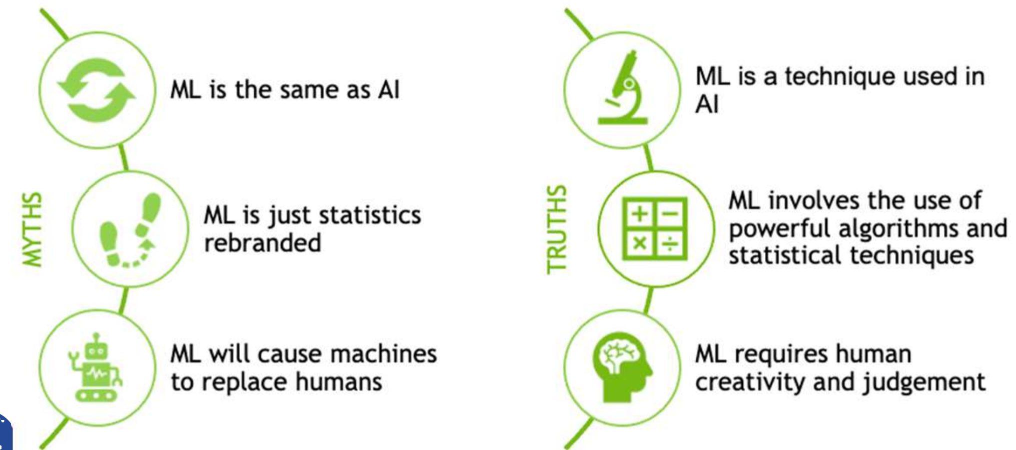

SPE International

Machine Learning

Webinars

ESSENCE OF MACHINE LEARNING

"A pattern exists…

We cannot pin it down mathematically…

We have data on it…"

Yaser Abu-Mostafa, Professor of EE & CS, Caltech

SPE International

Components of an ML Model

Webinars

Learning from data without being explicitly programmed

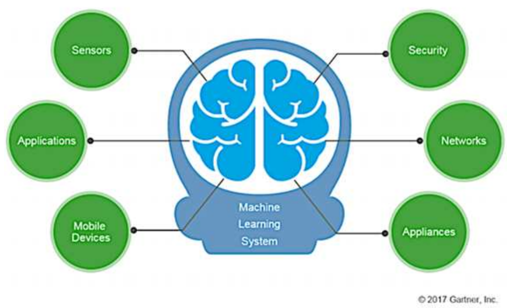

What is Required to Create Good ML Systems?

- Large Data Sets
- Data Processing Tools
- Powerful Algorithms
- Scalability
- Human Expertise

SPE International

Type of ML Algorithms

Webinars

TYPES	USE CASES / CHALLENGES
Supervised	Demand Forecasting, Recommender Systems
Unsupervised	Customer Segmentation, Fraud Detection
Reinforcement	Robotics, Process Optimization

SPE International

Supervised Learning Explained

Webinars

Example from Vending Machine: Coin Recognition

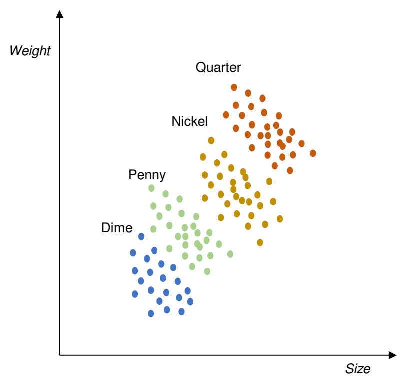

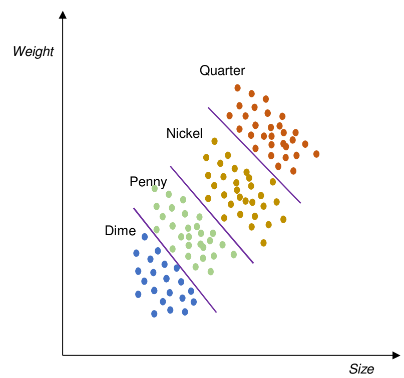

SPE International

Unsupervised Learning Explained

Webinars

Example from Same Vending Machine and Same Set of Coins

What it does:

- Finds hidden patterns & structures in unlabeled input data
- Determines correlations between data points
- Great for cluster analysis – exploratory data analysis to find hidden patterns or groupings in data

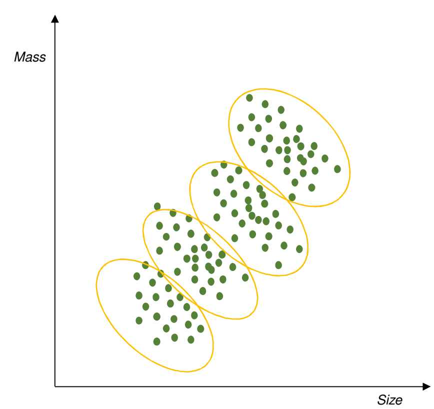

Instead of (input, correct output), we have (input, ?)

SPE International

Reinforcement Learning Explained

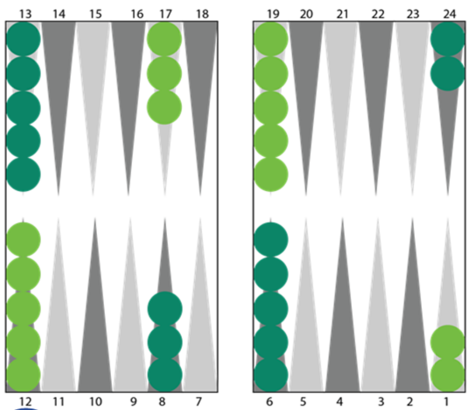

Instead of (input, correct output),

we have (input, *some* output, a grade for this output)

Reinforcement Learning

- process of autonomous learning through exploration
  - incremental reward-based progress (reinforcement)
  - dynamic programming in tandem with deep neural network approximators

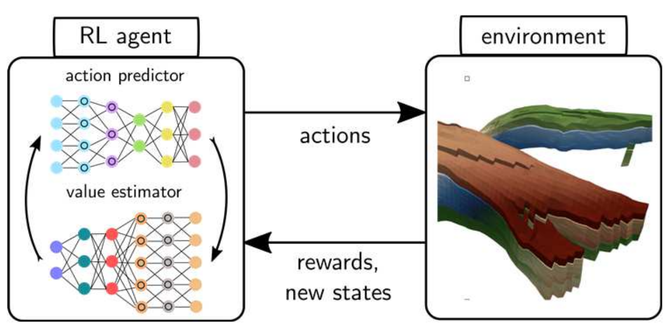

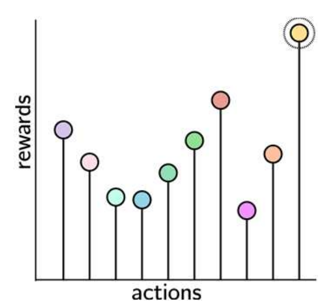

Growth of RL in HPC/AI

- exponential growth of AI/DL driven advancements in recent years
  - large investment of computational power driven by silicon revolution (GPUs, TPUs, ASICs)
- reinforcement learning impactfully features in multiple successful HPC/AI deployments
- focus on the application of HPC/AI specifically reinforcement learning in energy

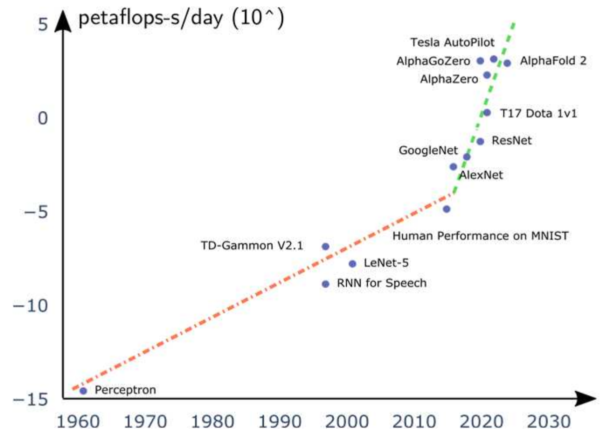

Reinforcement Learning Paradigms

- learning environment may be real (sensory inputs) or simulated
- single or multi-agent systems (adversarial or cooperative)
- may have varying degrees of application/model-dependance
  - model-based vs. model-free strategies
- highly suited for parallelism

Reinforcement Learning Paradigms

Webinars

- multiple approaches of reinforcement learning

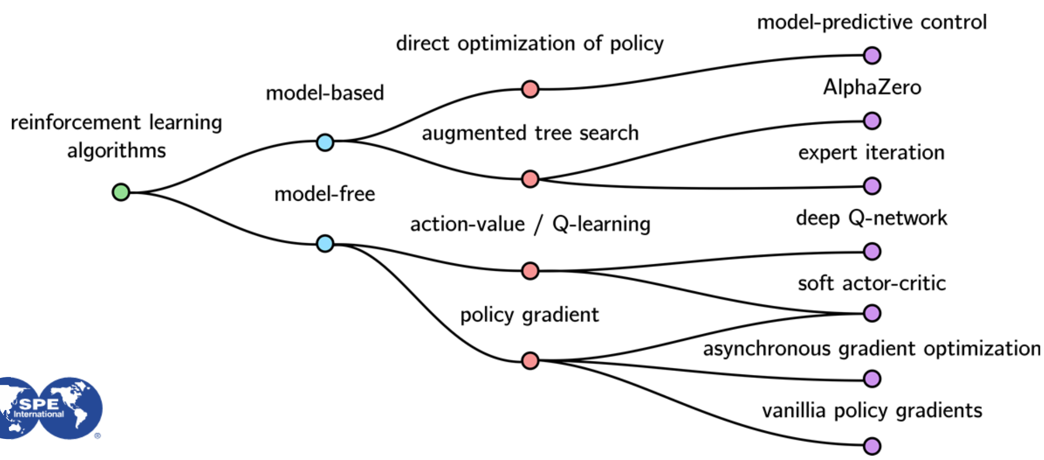

Role of Deep Learning in RL

Webinars

- deep neural networks function within the reinforcement learning frameworks:
  - surrogate models to enable faster emulation of environment
  - predicting best policy given state
  - estimating value given an action and state

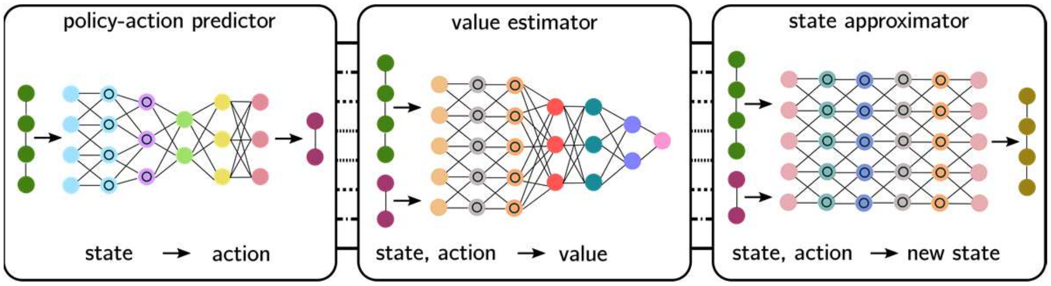

Monte Carlo Tree Search/RL

Webinars

inputs
- initial PUCT parameters
- configured training cycles
- recursive tree setup

environment step iteration

⇒ initiate training cycle, using value neural net
⇒ perform action maximizing PUCT parameter, based on neural net recommendation
⇒ regularize PUCT with Boltzmann temperature
⇒ observe new state, evaluate next child node
⇒ calculate reward once final leaf node reached
⇒ backpropagate reward up path to root node
⇒ if trained network beats previous cycle result:
   ⇒ retrain policy val network to minimize deviation
   ⇒ repeat training cycle

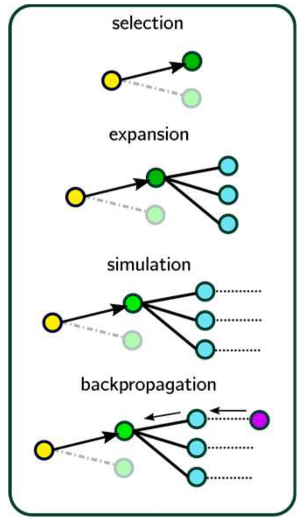

$$U(s,a) = \frac{W(s,a)}{N(s,a)} + C_{PUCT} \cdot ((1-\epsilon)P(s,a) + \epsilon \mathcal{D}(\sigma)) \cdot \frac{\sum N_c(s,a)}{1+N(s,a)}$$

PUCT (Predictor + Upper Confidence bounds for Trees)

SPE International

Soft-Actor-Critic Algorithm

Webinars

inputs
- initial policy parameters
- Q-targets
- replay buffer setup

environment step iteration

⇒ perform predicted action with perturbation
⇒ observe new state
⇒ calculate reward
⇒ store state-action pairs, rewards, life of environment within replay buffer
⇒ sample batch for ML training
⇒ *if state update is performed:*
   ⇒ update MSBE (mean-square Bellman loss)
   ⇒ update two Q-estimates through MSBE and stochastic actions
   ⇒ use minimum of Q-estimates & KL-divergence to perform actor updates
   ⇒ update coefficient of entropy regularization
   ⇒ update Q-targets

SPE International

HPC Implementation & Parallelism

- performing training and inference on dedicated GPUs
- minimize memory copies from device-host & device-device
- simultaneous inference during parallel self-play
  - our models utilize MPS (multi-process service) capability
  - MiG capability in A100 GPUs
- reconfiguration/reallocation and optimization of GPU resources
  - bandwidth vs. compute dominated phases (inference, training, sim)
- usage of mixed precision when accuracy is not necessary

Path Planning

- problem task: finding non-self-intersecting minimum energy path between two points in a 3D landscape
- real time robotic navigation given a topographical map
- state: current path and position
- action: move
- reward: sum of total energy during arc traversal

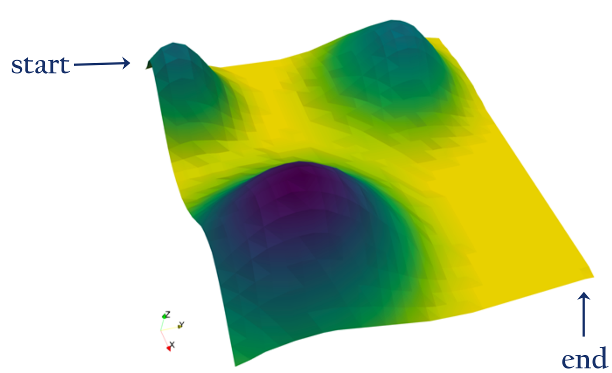

Path Planning

- progress during training cycles:

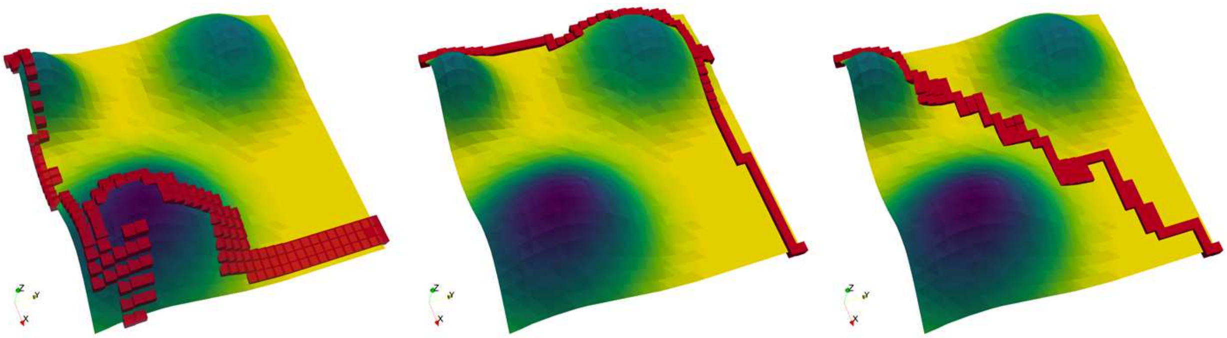

Open Pit Optimization

- mine pit optimization is an important problem in the mining industry
- avoiding slope instability requires geometric and temporal optimization of mining
- maximize NPV (net present value) of extractable ore
- state: mined area/time, action: next set of mining points, reward: cumulative ore NPV

Open Pit Mine, Hexagon Corp.

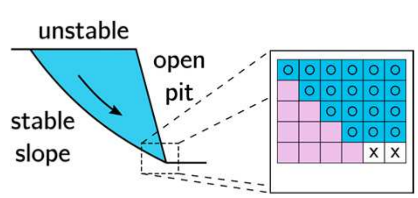

precedence arcs need to be extracted first

Open Pit Optimization

Webinars

- discretize ore/mining value
  - temporal discounting factor
  - limited operations per time period for mining

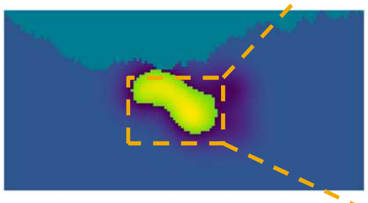

-22.0	-39.0	-22.0	-33.0	-19.0	-28.0	-30.0	12.0	-10.0	600.0	-349.0	-30.0	-37.0	-21.0	-13.0	650.0
-46.0	-51.0	-59.0	-49.0	-65.0	412.0	842.0	639.0	-44.0	550.0	332.0	-47.0	-62.0	594.0	-62.0	-48.0
-68.0	-86.0	-100.0	-80.0	-89.0	-43.0	-73.0	-87.0	230.0	-77.0	-67.0	-76.0	167.0	526.0	-85.0	-83.0
-107.0	-99.0	-120.0	-106.0	-94.0	684.0	-81.0	-98.0	383.0	-99.0	601.0	128.0	9600.0	-89.0	-85.0	-112.0
-110.0	-124.0	-120.0	-142.0	-28.0	550.0	-123.0	-138.0	414.0	-120.0	452.0	-124.0	-140.0	-140.0	-125.0	-108.0
-151.0	-147.0	-132.0	-145.0	-50.0	570.0	-139.0	-147.0	362.0	264.0	6820.0	575.0	388.0	-164.0	-136.0	-147.0
-188.0	-183.0	-175.0	322.0	351.0	-175.0	-181.0	741.0	475.0	334.0	269.0	-180.0	-174.0	-168.0	-172.0	-182.0
-198.0	-208.0	-180.0	240.0	250.0	-202.0	-214.0	406.0	557.0	-205.0	-189.0	-194.0	-215.0	-184.0	-198.0	-203.0

(a) iteration = 0, value = 8893

-22.0	-39.0	-22.0	-33.0	-19.0	-28.0	-30.0	12.0	-10.0	600.0	-349.0	-30.0	-37.0	-21.0	-13.0	650.0
-46.0	-51.0	-59.0	-49.0	-65.0	412.0	842.0	639.0	-44.0	550.0	332.0	-47.0	-62.0	594.0	-62.0	-48.0
-68.0	-86.0	-100.0	-80.0	-89.0	-43.0	-73.0	-87.0	230.0	-77.0	-67.0	-76.0	167.0	526.0	-85.0	-83.0
-107.0	-99.0	-120.0	-106.0	-94.0	684.0	-81.0	-98.0	383.0	-99.0	601.0	128.0	9600.0	-89.0	-85.0	-112.0
-110.0	-124.0	-120.0	-142.0	-28.0	550.0	-123.0	-138.0	414.0	-120.0	452.0	-124.0	-140.0	-140.0	-125.0	-108.0
-151.0	-147.0	-132.0	-145.0	-50.0	570.0	-139.0	-147.0	362.0	264.0	6820.0	575.0	388.0	-164.0	-136.0	-147.0
-188.0	-183.0	-175.0	322.0	351.0	-175.0	-181.0	741.0	475.0	334.0	269.0	-180.0	-174.0	-168.0	-172.0	-182.0
-198.0	-208.0	-180.0	240.0	250.0	-202.0	-214.0	406.0	557.0	-205.0	-189.0	-194.0	-215.0	-184.0	-198.0	-203.0

(b) iteration = 1840, value = 14327

SPE International

Field Development Planning

Webinars

- maximization of NPV (net present value) through spatiotemporal well placement in a benchmark model (Egg reservoir, Jansen et. al. 2014)
- state: saturations/pressures and previously drilled well locations
- actions: new well locations at given timestep
- reward: cumulative NPV

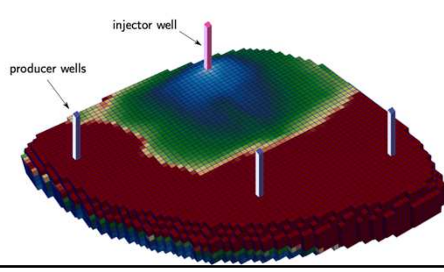

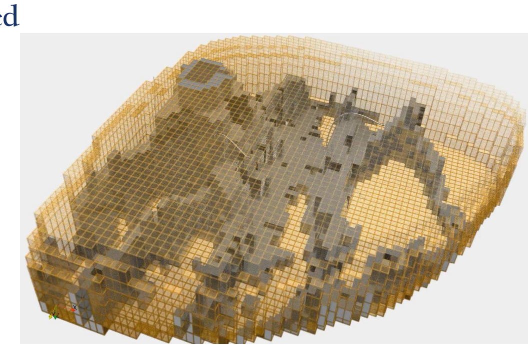

SPE International

Field Planning Optimization

- hybrid run with 90 CPU-cores, 2 A100 GPUs for vertical wells only (2D optimization)
- 165,667 action-exploration forward reservoir simulations within 36 hours

net present value

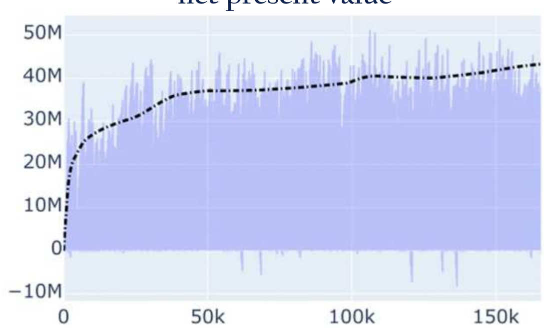

Field Planning Performance Benchmark

- benchmarking performed with details of intra/inter-operation performance
- metrics chosen to compare performance benchmark
  - best solution found in NPV (vertical wells), normalized NPV (trajectories)
  - distribution of NPV solutions indicating power of method
  - time taken to reach episodes as well as speed instantaneously
- series of calculations with fixed time (120k seconds) runs for 20, 96 & A100 runs:
  - ratio of number of actions inferred at 120k seconds
  - speed comparison and peak performance difference

Field Planning Results

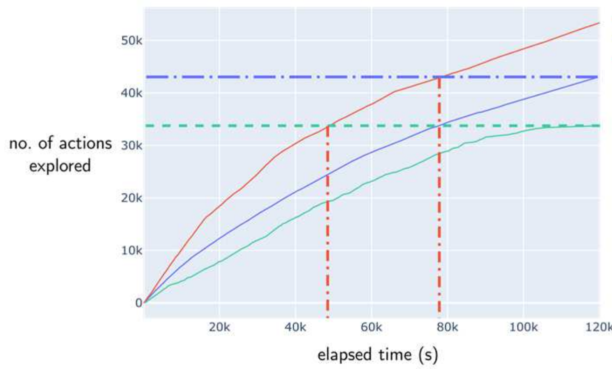

normalized speed, actions inferred per second

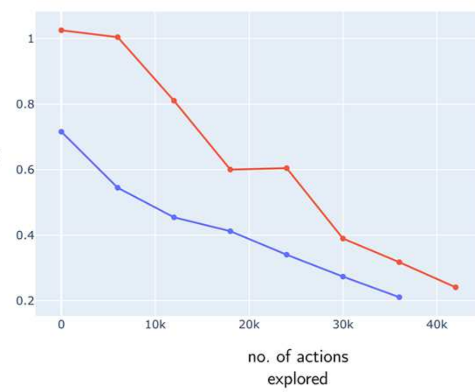

HPC Benchmarks – Field Planning

metric (a) time elapsed to reach fixed explorations	
A100 over 20-core	A100 over 96-core
245.4%	152.9%

metric (b) speed of reinforcement learning	
A100 over 20-core	A100 over 96-core
252.8% (mean) 1169.5% (peak)	147.9% (mean), 184.3% (peak)

metric (c) NPV normalized by length	
A100 over 20-core	A100 over 96-core
35.2%	7.83%

Conclusions

Webinars

- optimization based on reinforcement learning shows promise in a range of scenarios
- both model-based and model-free DRL strategies have industrial applications
- GPUs enable HPC/AI deployments
  - massive parallelism in inference & training
  - ease of deployment in the cloud
- exploring similar approaches coupled to high-fidelity simulations
  - carbon capture, power plant optimization and other clean energy initiatives

SPE International