

Upstream Analytics

Upstream Analytics

Industry Day

February 25-26, 2020

Machine Learning
at
Sandia



Stephen J. Verzi

Outline




- Predicting Permeability in Porous Media using Physics-informed Deep Neural Networks (DNNs)
- Adaptive Optimal Control of Stable Grid State Space Navigation with Reinforcement Learning (RL)
- Anomaly Detection with Deep Spiking Neural Networks (SNNs)
- Machine Learning (ML) Basics and Potential Uses

Predicting Permeability in Porous Media Using DNNs



Synthetically generated data

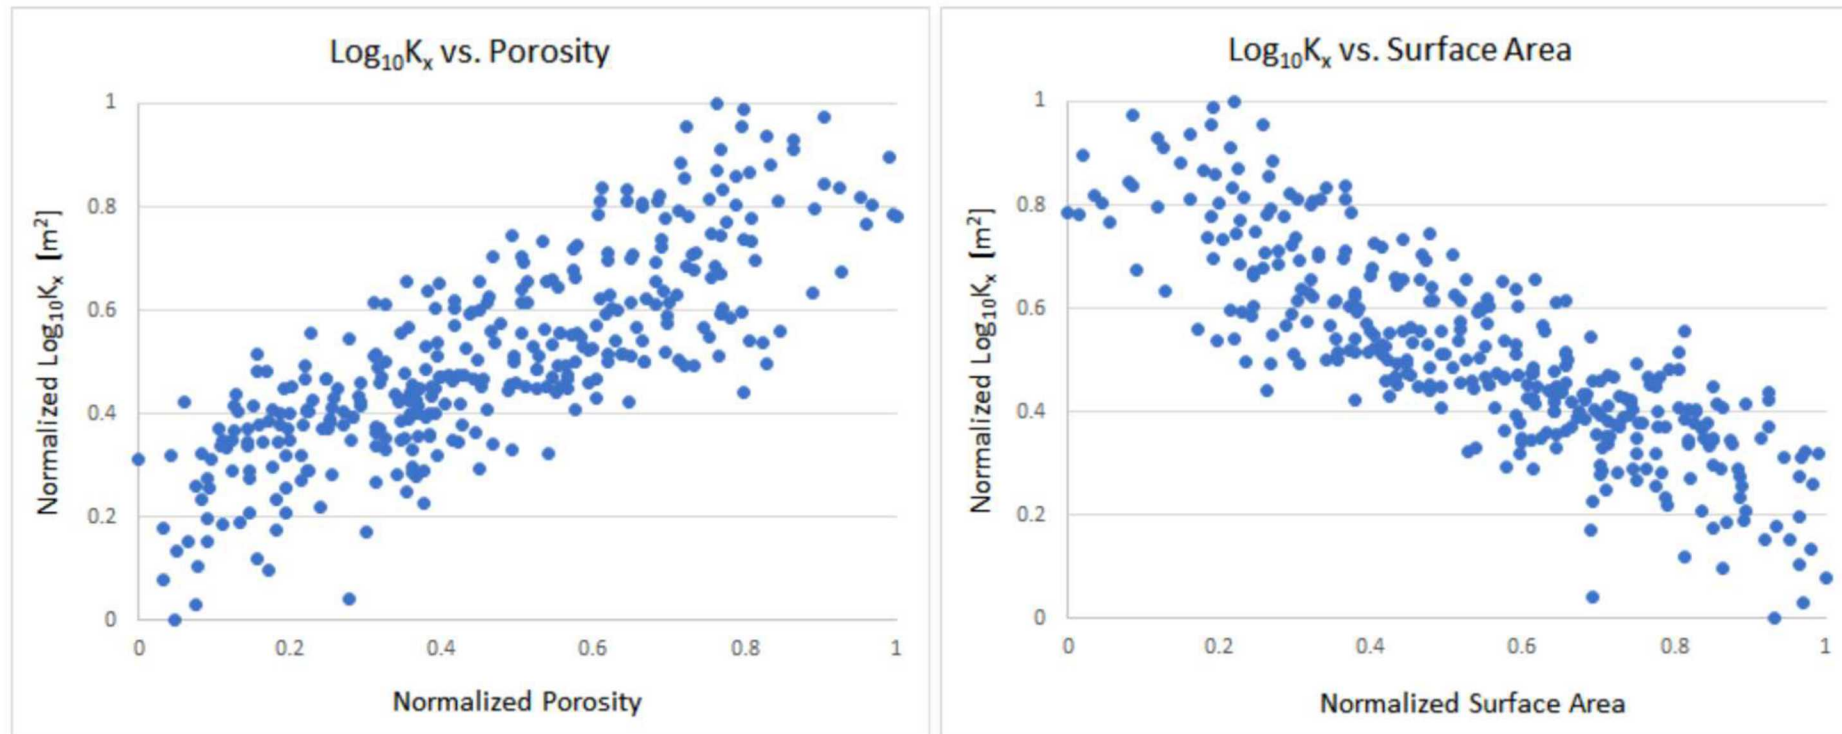


Examples of synthetic porous media generated with different sizes of sphere. Fluid flows through the void space in black. Permeability decreases from left to right and ranges over two orders of magnitude in m^2 .

Predicting Permeability in Porous Media Using DNNs



Correlation with porosity and surface area

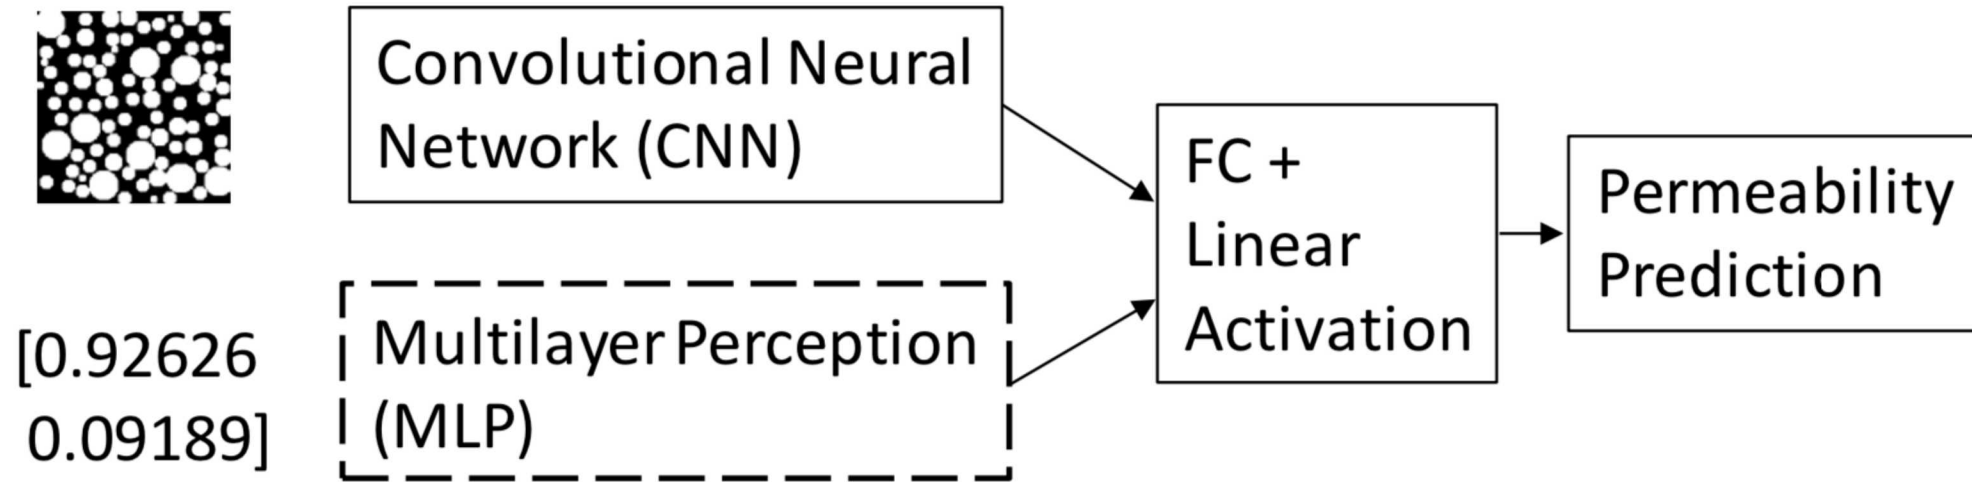


Normalized permeability in x-direction ($\log_{10}K_x$) vs. porosity (left) and surface area (right).

Predicting Permeability in Porous Media Using DNNs



Physics-informed machine learning – using porosity and surface area




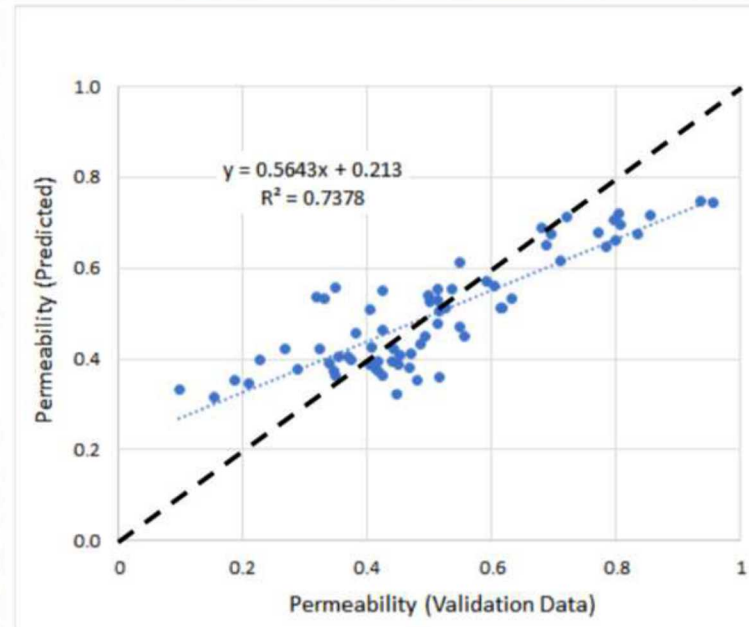
Schematic of physics-informed machine learning architecture.

Predicting Permeability in Porous Media Using DNNs



Results

Model	Final Average Validation MSE
CNN-only	0.011
Physics-informed architecture	0.007



Comparison of permeability prediction using validation data for CNN-only (left) and physics-informed (right) architectures (linear regression fit with blue dotted line and theoretically perfect prediction with the black dashed lines are also shown).

Stable Grid State Space Navigation with RL



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.

Figure from NERC Reliability Concepts V 1.0.2


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

Stable Grid State Space Navigation with RL



“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



Transient Stability Level Curves

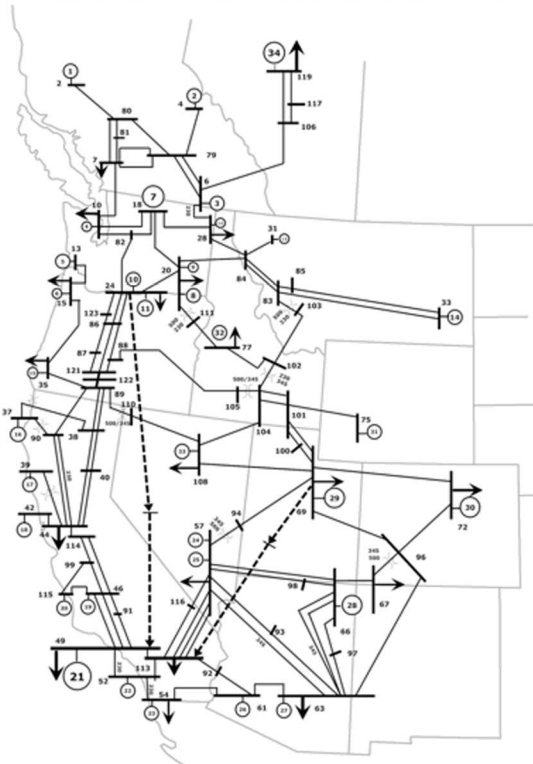


Transient & Voltage Stability Level Curves

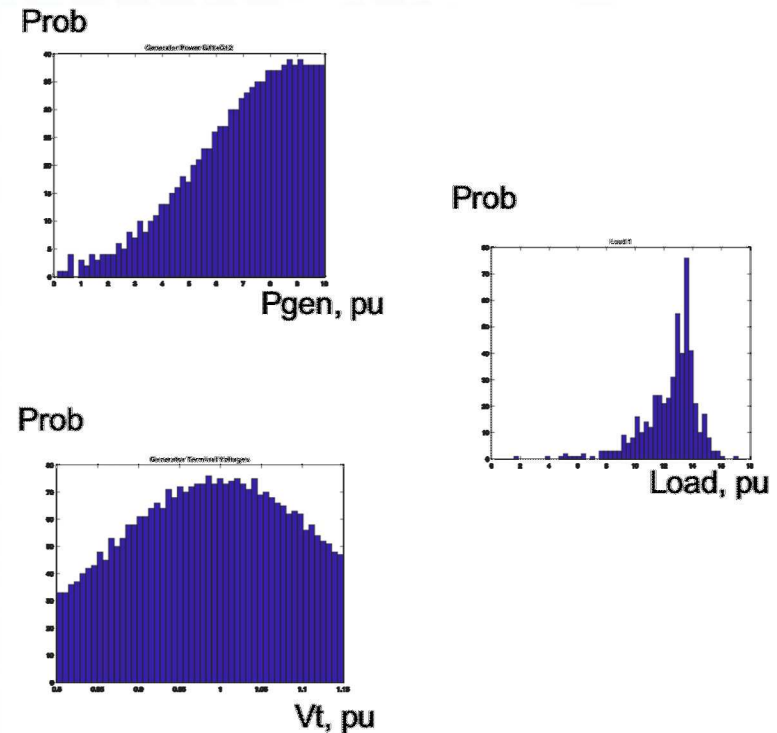
Stable Grid State Space Navigation with RL



Defining the Grid State Space



miniWECC reduced order model [Neely, Johnson, Byrne & Elliot, 2015]



Grid State Space Variables/Distributions

Stable Grid State Space Navigation with RL



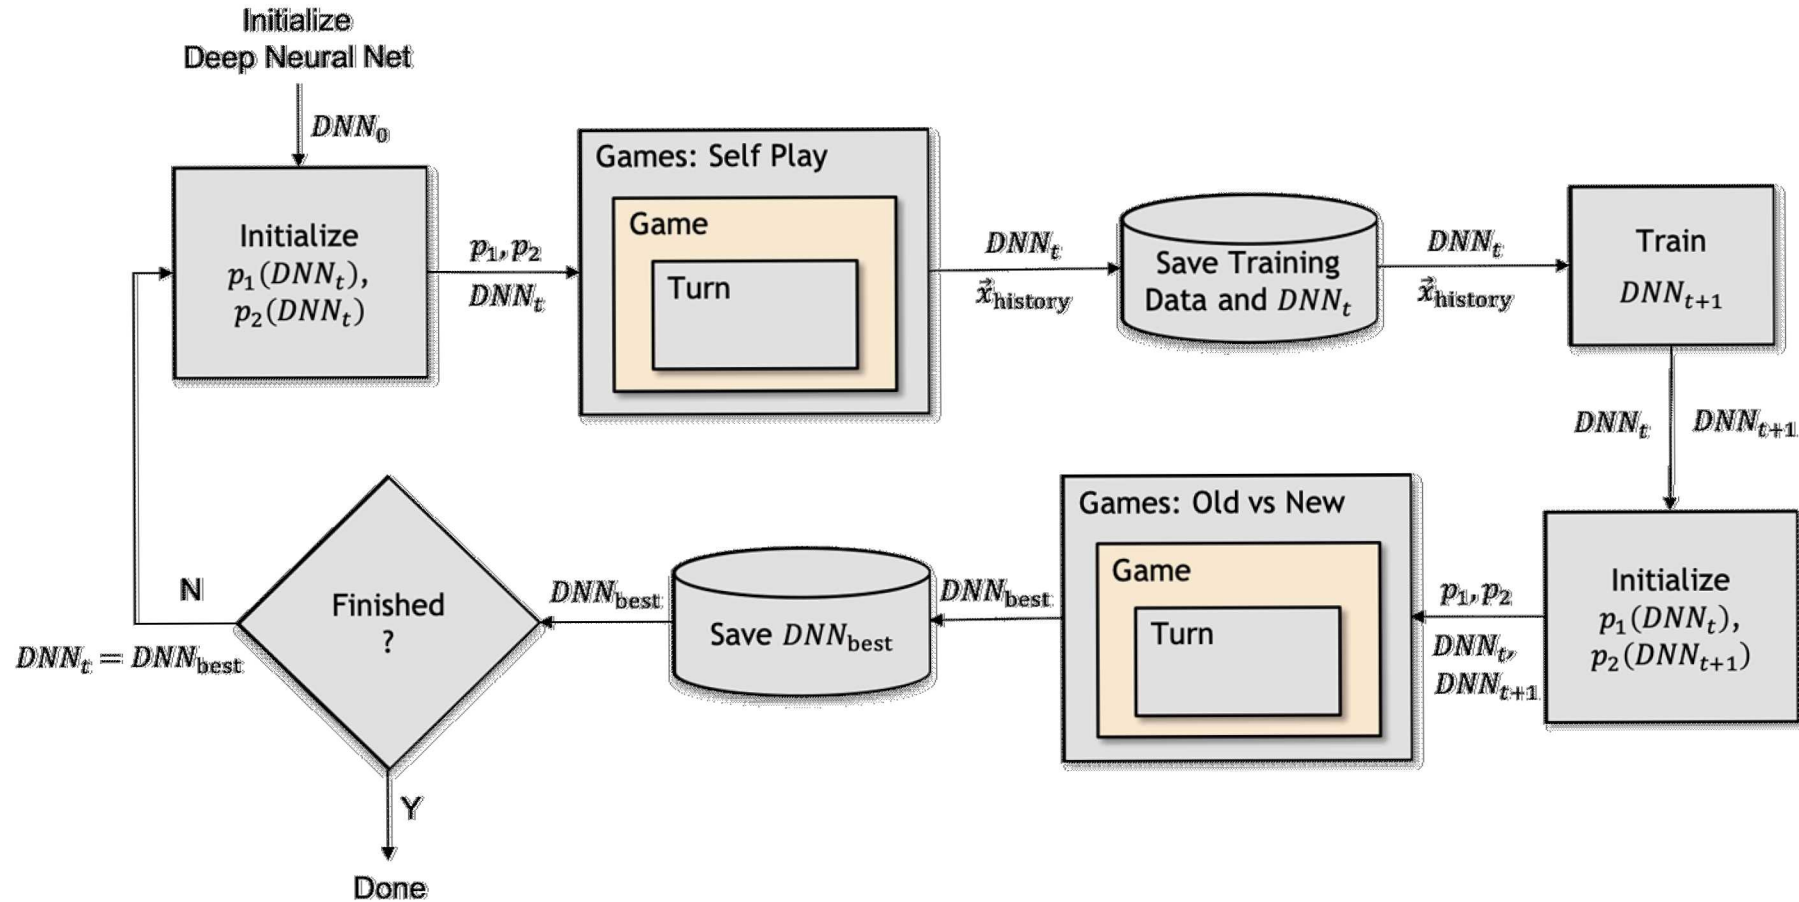
The AlphaGrid Game:

- **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

Stable Grid State Space Navigation with RL



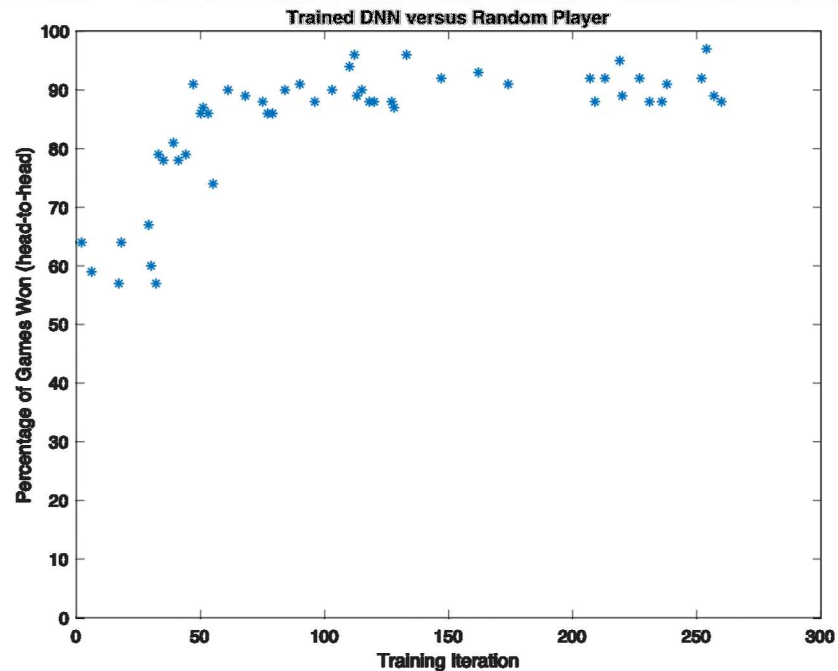
RL – DNN Block Diagram



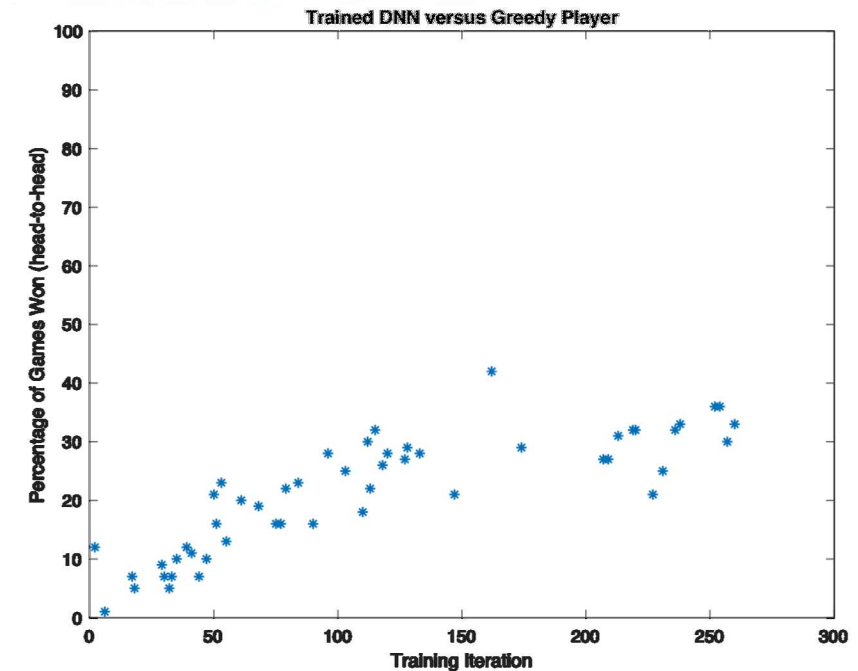
Stable Grid State Space Navigation with RL



Results



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



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

Anomaly Detection with Deep (SNNs)




Phase-coding using spiking neurons

Discretized leaky integrate and fire (LIF)



- $u_j(t) = (u_j(t-1) - \lambda_j(u_j(t-1) - u^{eq})) (1 - z_j(t-1)) + u^{eq} z_j(t-1) + \sum_i w_{ij} x_i(t) + \sum_q w_{qj} \rho_q(z_q(t-1))$
- if $u_j(t) \geq \theta_j$ and $u_j(t-1) > 0$ then the neuron spikes,

where


- u_j is the neuron potential
- λ_j is the rate of leakage
- u^{eq} is the resting potential
- $w_{ij} x_i(t)$ is weighted current injection
- θ_j is the threshold
- $w_{qj} \rho_q(z_q(t-1))$ are weighted phase-coded spikes



ANN



SNN




Anomaly Detection with Deep (SNNs)



Phase-coding using spiking neurons

Define phase-coding as

- $$\rho_j(z_j(t)) = z_j(t) \left(\left((k-1) - t \bmod \frac{k}{\tau_{\text{step}}} \right) - \rho_j^0 \right)$$
- within a window of k steps each of length τ_{step} in the history
- ρ_j^0 is the phase-code reference value (see Time-encoding)
- actual spike time is $t + \frac{(k-1) - t \bmod k}{k}$ where $\tau_{\text{step}} = 1$ at

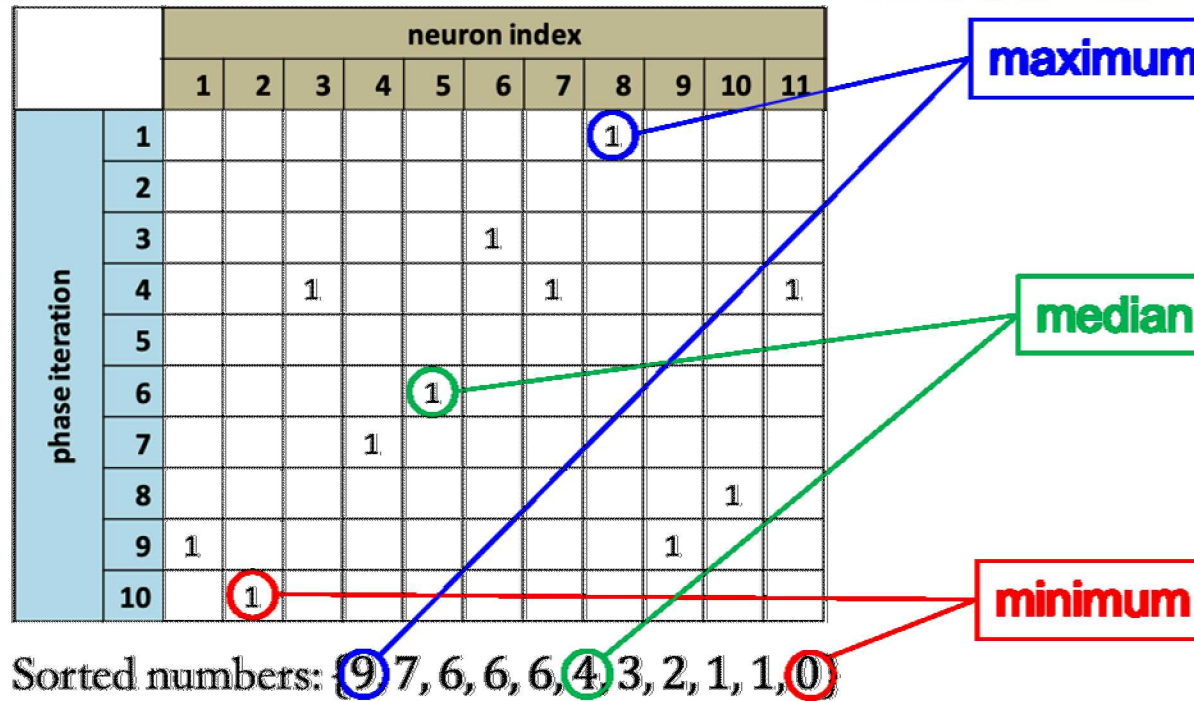


Anomaly Detection with Deep (SNNs)



Sorting with spiking neurons (spiking-sort)

- input = {1, 0, 6, 3, 4, 7, 6, 9, 1, 2, 6} with $k = 10$ [Verzi et al., 2018]





Anomaly Detection with Deep (SNNs)



Phase-coding demonstration

- Caltech 101 soccer ball image [Fei-Fei et al., 2004]



Anomaly Detection with Deep (SNNs)



Median filtering with spiking neurons

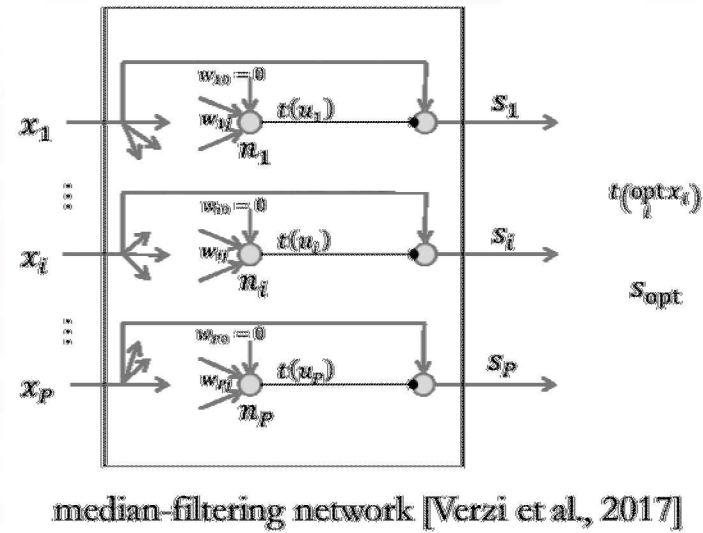
Using optimization-based, signed-rank formulation for $\{x_1, x_2, \dots, x_N\}$

- $\tilde{R}(x) = \sum_{i=1}^N \text{sign}(x - x_i)$
- the median, \tilde{x} , is such that $\tilde{R}(\tilde{x})$ is closest to 0, also $-\frac{N-1}{2} \leq \tilde{R}(x) \leq \frac{N-1}{2}, \forall x$
- we apply median-filtering to each pixel and its neighbors

x_1	x_2	x_3
x_4	x_5	x_6
x_7	x_8	x_9



image from Cal Tech 101 [Fei-Fei et al., 2004]



Anomaly Detection with Deep (SNNs)



Spiking Adaptive Median Filtering (AMF)

Neural-inspired simplification of adaptive center-weighted median-filtering (ACWMF) [Chen and Wu, 2001]

$$o_{ij} = \begin{cases} \hat{\rho}_{ij}^1, & \text{if } \exists m, \hat{\rho}_{ij}^m > \theta^m \\ x_{ij}, & \text{otherwise} \end{cases}$$


Here $m \in \{1, 2, 3, 4\}$ and $s = 3$

$$\hat{\rho}_{ij}^m = \text{median}\{x\}_{x \in \Omega_{ij}^m}$$

$$\Omega_{ij}^m = \{x_{lr} \mid i - m \leq l \leq i + m, j - m \leq r \leq j + m\}$$

$$\theta^m = s \cdot \text{median}\{\hat{\rho}_{ij}^m - x\}_{x \in \Omega_{ij}^m} + \delta^m$$

$$\delta^m = \frac{((2m+1)^2 - 1)}{2}$$

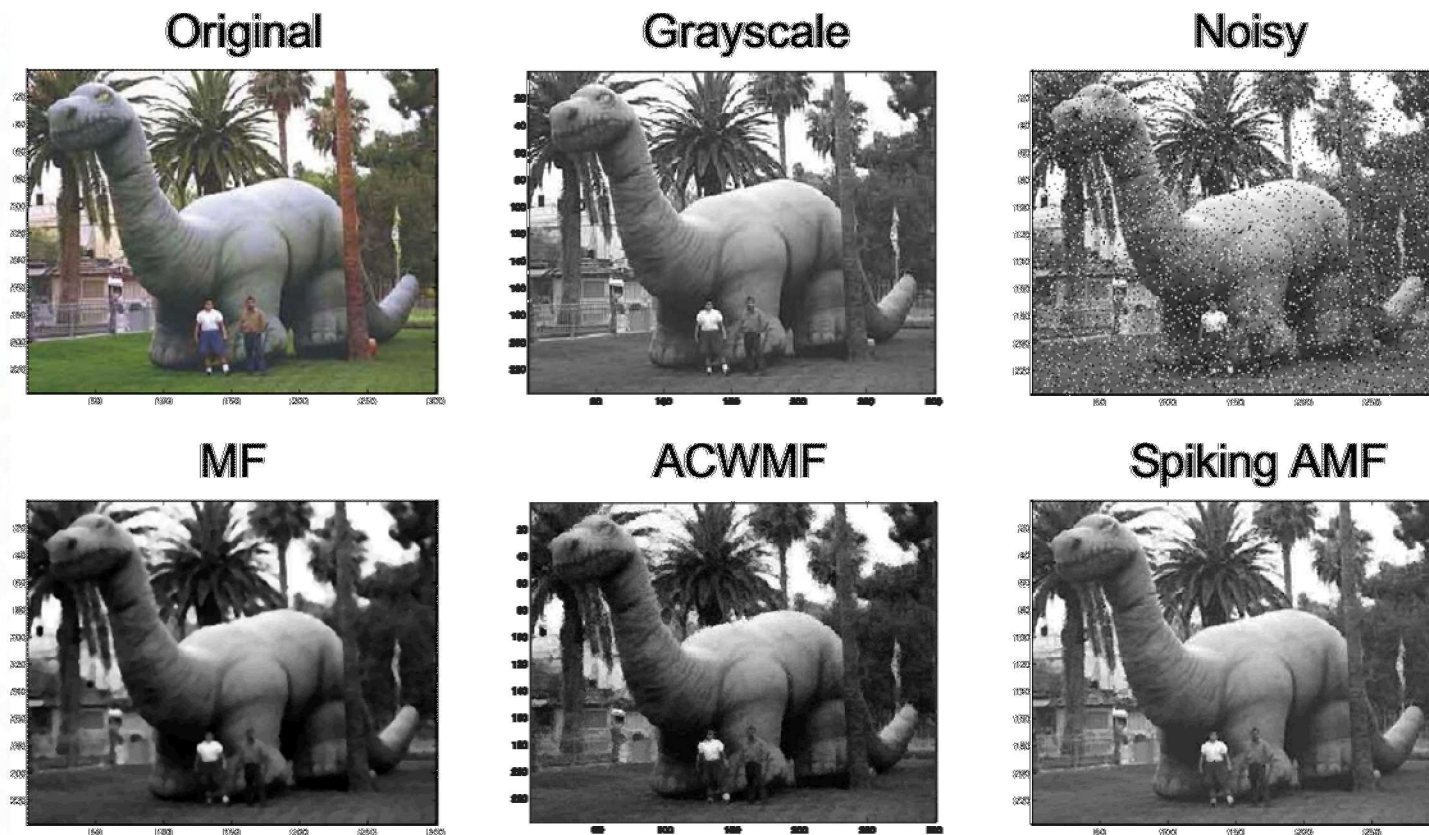


Anomaly Detection with Deep (SNNs)



Results

- Validate using random noise with Cal Tech 101 images [Fei-Fei et al., 2004]



Anomaly Detection with Deep (SNNs)



Results

Noise type	percent of pixels changed				average pixel value changed (changed pixels only)			
	Noisy image	MF	ACWMF	Spiking AMF	Noisy image	MF	ACWMF	Spiking AMF
10% “pepper”	9.55	66.86	25.46	12.69	12.6 (131.6)	9.6 (13.7)	6.7 (23.6)	4.2 (29.0)
10% “salt”	9.64	66.87	25.57	12.84	12.9 (134.7)	9.7 (13.8)	6.8 (24.0)	4.1 (28.0)
10% “salt & pepper”	9.60	66.29	25.03	12.54	12.8 (134.5)	9.4 (13.4)	6.5 (23.0)	3.9 (27.3)
20% “random impulse”	19.94	69.27	31.83	19.98	17.4 (87.0)	10.3 (14.3)	7.8 (22.5)	5.2 (24.0)

Using 12 different images and 10 random seeds for each type of noise

ML Basics and Potential Uses



Definition

- Perform “better” next time
- Performance improves over time

Examples

- Regression
- Classification

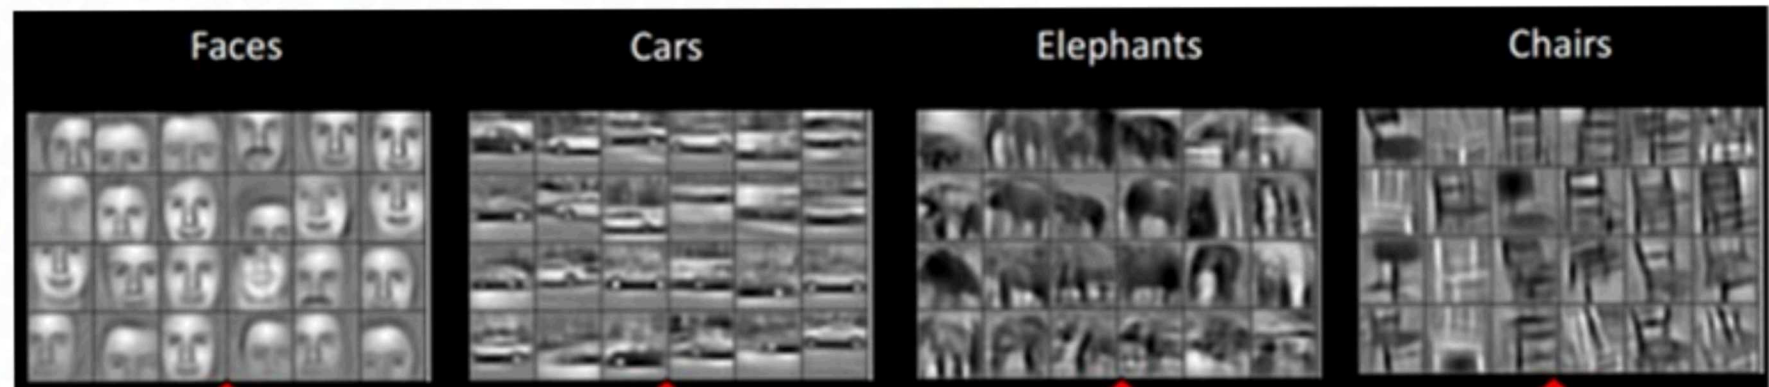







image from [Siegel et al., 2016]

ML Basics and Potential Uses




Deep Learning (DL)

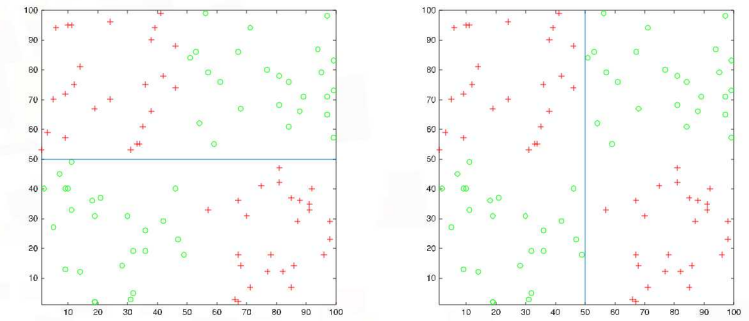
- Use multilayer perceptron (MLP) to implement ML
- Perceptron




- 1 hidden layer



perceptron decision boundary:
hyperplane



multi-layer perceptron decision boundary:
continuous function




ML Basics and Potential Uses




Function approximation


- ML
 - Continuous function f can be approximated with simple functions s_i
 - $f(x) \approx \sum_i c_i s_i(x)$
 - ML can be used to find c_i through training
- DL
 - An MLP represents s_i with its combined neural activation functions



- 1 hidden layer



- Multiple hidden layers



ML Basics and Potential Uses



Potential uses

- Exact equations known – use them not ML
- Parameterized functional form known – can use ML to find parameter values
- Functional form not known – can use DL to find an approximate functional mapping, i.e., supervised learning
- For control and/or dynamic system interaction – can use RL for adaptive optimal control

Selected References



Chen, T., Wu, H.R.: Adaptive impulse detection using center-weighted median filters. IEEE Signal Processing Letters 8(1), pp. 1–3. IEEE, 2001. DOI 10.1109/97.889633. URL <https://dx.doi.org/10.1109/97.889633>

Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: an incremental Bayesian approach tested on 101 object categories. IEEE. CVPR 2004, Workshop on Generative-Model Based Vision. 2004.

Guttromson, R., Verzi, S.J., Lawson, L., Levin, D., Melander, D., Sorensen, A., Cauthen, K., Wilches-Bernal, F., Berg, T., Lavrova, O. Integrated Cyber/Physical Resiliency Modeling. Sandia Report, SAND2018-876234. 2018.

Hill, A. J., Rothganger, F., Vineyard, C. M., Follett, D. R., Follett, P. L., Smith, M. R., Verzi, S. J., Severa, W., Wang, F., Aimone, J. B., Naegle, J. H., and James, C. D. (2017). A spike-timing neuromorphic architecture. In Proceedings of the International Conference on Rebooting Computing (ICRC), pp. 1–8. IEEE, 2017. DOI 10.1109/ICRC.2017.8123631. URL <https://doi.org/10.1109/ICRC.2017.8123631>

Hopfield, J. J. Pattern recognition computation using action potential timing for stimulus representation. Nature, 376:33–36. Nature, 1995.

Neely, J., Johnson, J., Byrne, R., Elliot, R. Structured optimization for parameter selection of frequency-watt grid support functions for wide-area damping. International Journal of Distributed Energy Resources and Smart Grid, 11 (1), pp. 69-94. 2015.

North American Electric Reliability Corporation (NERC), Reliability Concepts, Version 1.0.2.

Selected References



Siegel, C., Daily, J., Vishnu, A. Adaptive neuron apoptosis for accelerating deep learning on large scale systems. 2016 IEEE International Conference on Big Data (Big Data), pp. 753-762. 2016.

Verzi, S.J., Vineyard, C.M., Vugrin, E.D., Galiardi, M., James, C.D., Aimone, J.B.: Optimization-based computation with spiking neurons. In Proceedings of the IEEE 2017 International Joint Conference on Neural Network (IJCNN), pp. 2015–2022. IEEE, 2017. DOI 10.1109/IJCNN.2017.7966098. URL <https://dx.doi.org/10.1109/IJCNN.2017.7966098>

Verzi, S.J., Rothganger, F., Parekh, O.D., Quach, T.T., Miner, N.E., James, C.D., Aimone, J.B.: Computing with spikes: The advantage of fine-grained timing. Neural Computation (accepted), 2018.

Yoon, H., Melander, D., Verzi, S.J., Permeability Prediction of Porous Media using Convolutional Neural Networks with physical properties. In Proceedings of the AAI-MLPS 2020 Symposium on Combining Artificial Intelligence and Machine Learning with Physics Sciences. Accepted for presentation at AAI-MLPS, 2020.

Thank You



Questions?