



Exceptional service in the national interest

# Assessing a Neuromorphic Platform for use in Scientific Stochastic Sampling

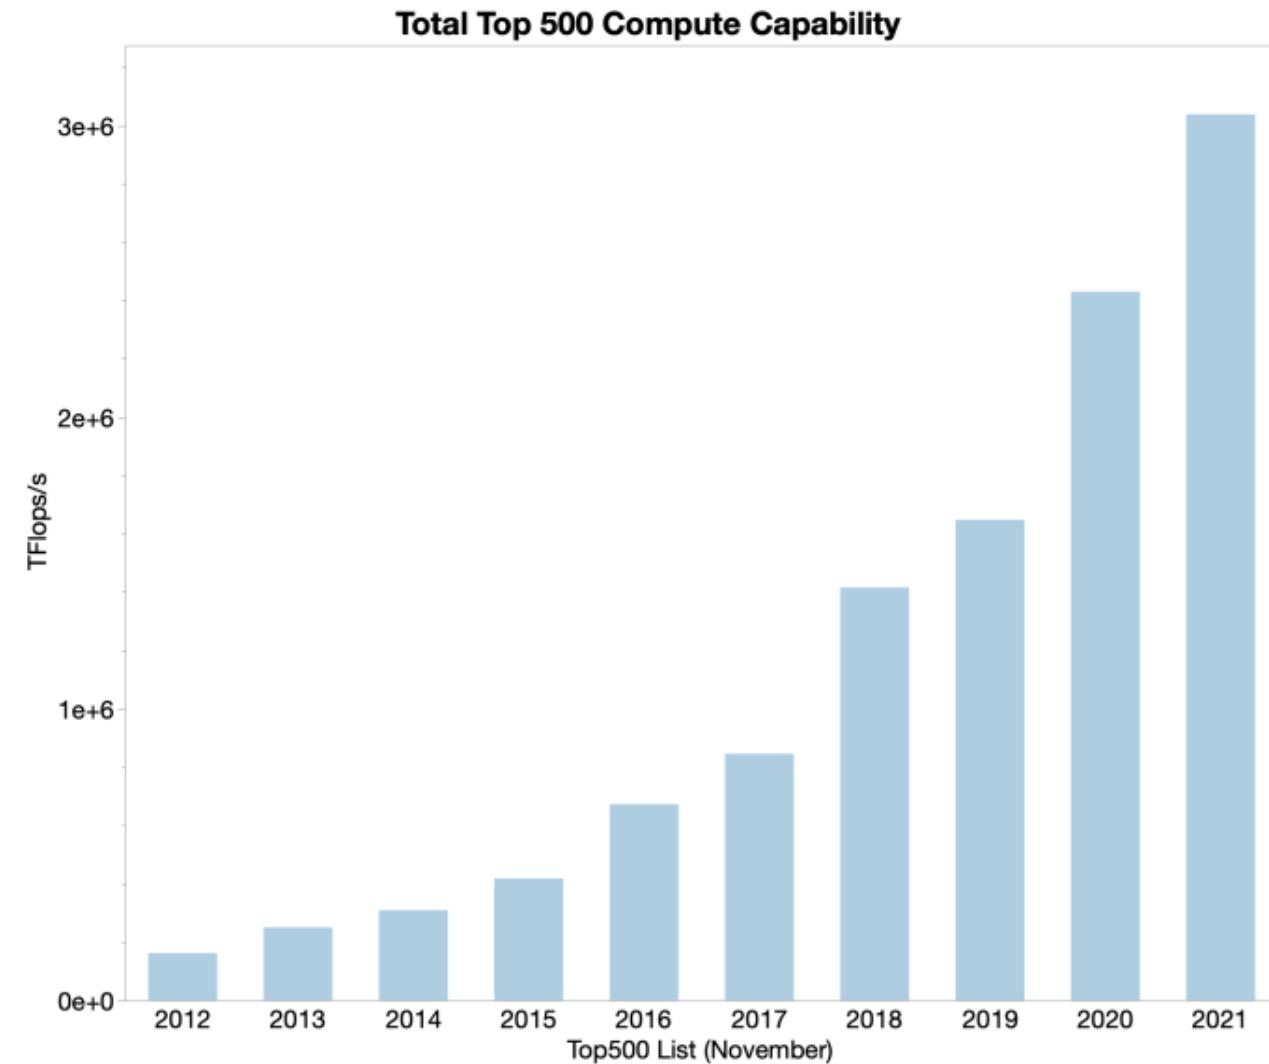
Darby Smith

May 12<sup>th</sup>, 2022

International Symposium on Roadmapping Devices & Systems

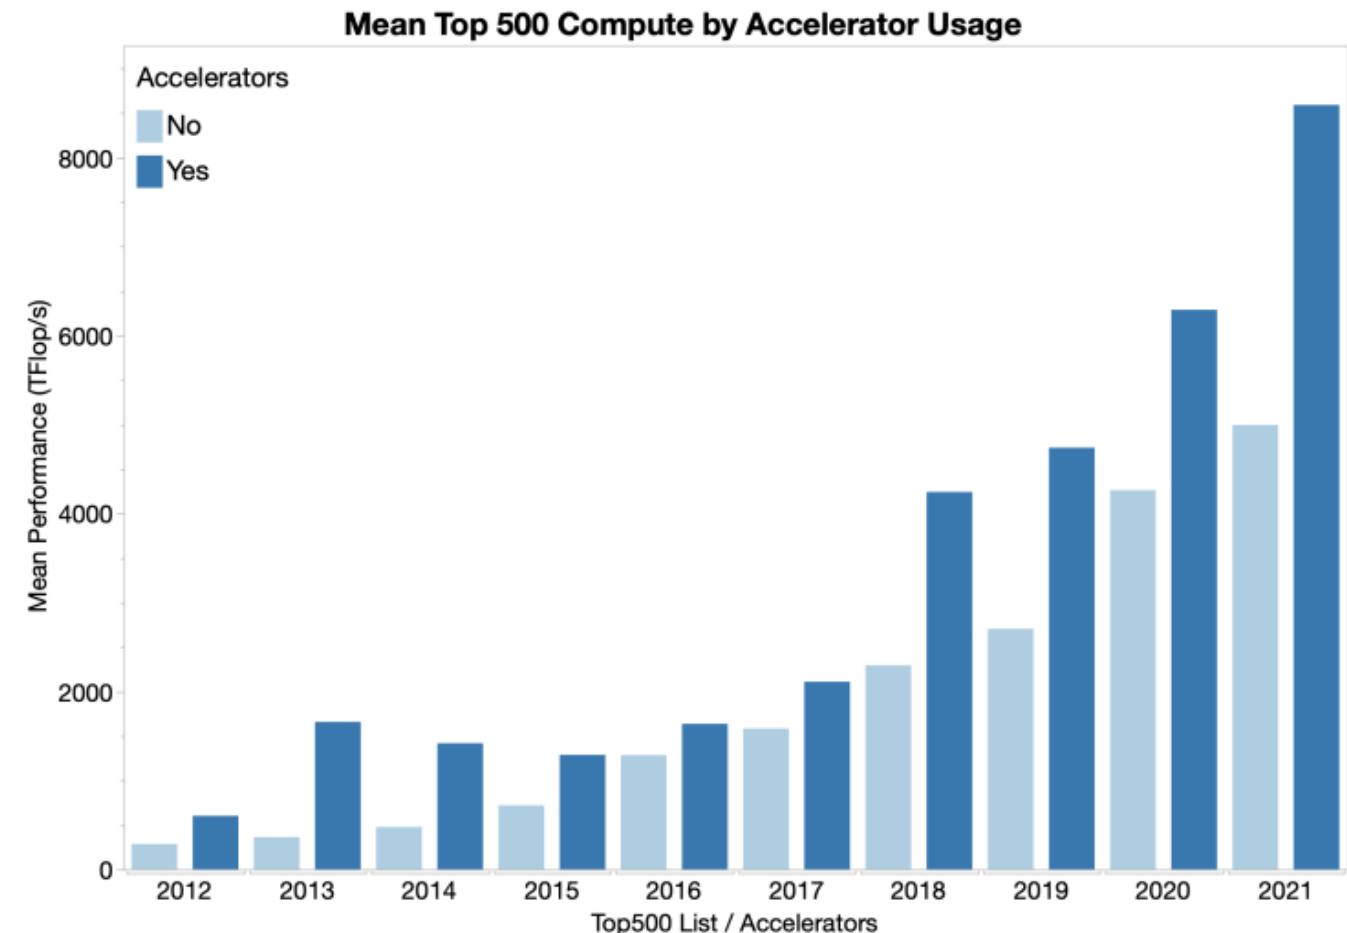
# Supercomputing – More Super Than Ever

- Computational loads are more demanding than ever
- Many reports of ending scaling principles
  - Moore's Law
  - Denard Scaling
- But, High Performance Computing is growing to meet this need regardless!
- No clear slowdown in large scale systems



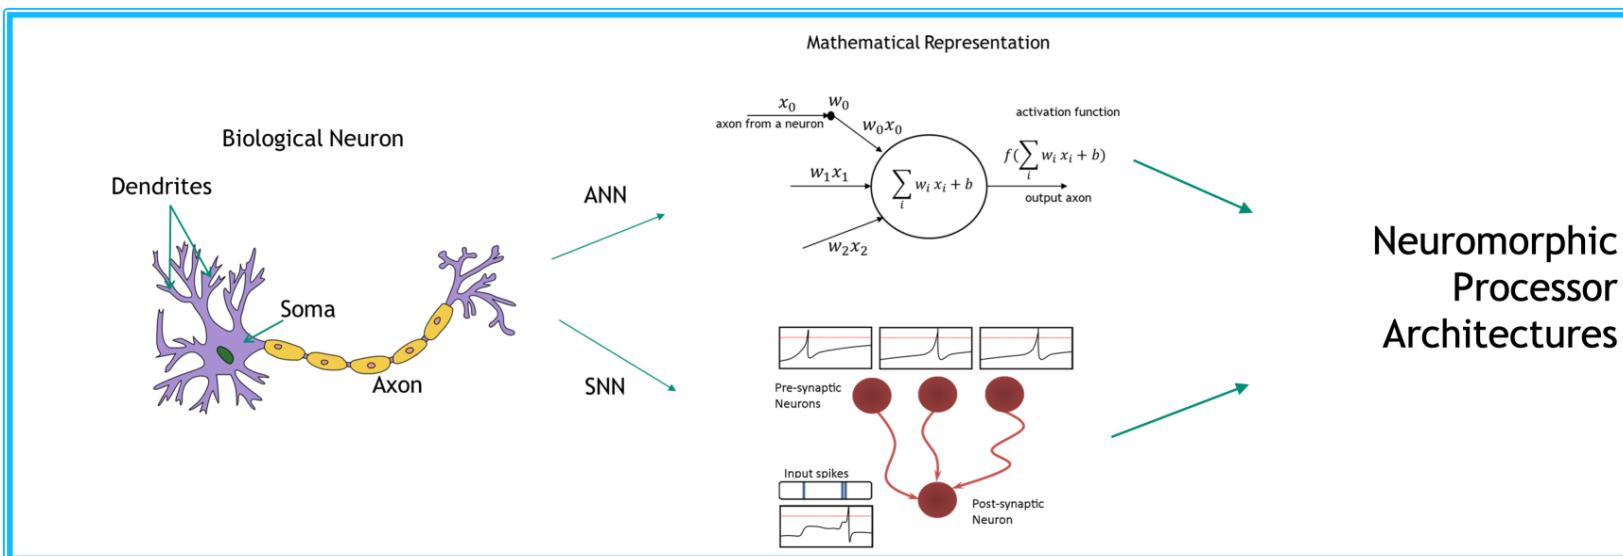
# Heterogenous Processors Enable Big Compute

- A growing percentage of large-scale compute relies on heterogenous accelerators
  - GPUs
  - Intel Phi
  - Systolic Arrays
  - Deep Learning Accelerators
- Heterogenous compute can improve computational density and **energy efficiency**



# Neuromorphic Platforms Offer Energy Benefits

- Neuromorphic processors mimic event-driven processing seen in biological neural networks
- Inherently parallel, these platforms can have millions of copies of 'neurons' at very lower power
  - But, each neuron has limited behavior
- Configuring a group of neurons to do something *useful* is the challenge



# Neuromorphic's Potential Spans Fields

## *Scientific Computing*

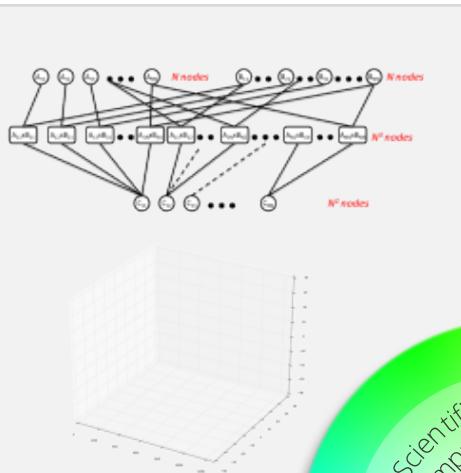
Well-understood requirements

Opportunity:

*Novel neuromorphic algorithms*

Examples:

- Solving SDEs (Monte Carlo PDE solutions)
- Neural Graph Analytics



## *Machine Learning*

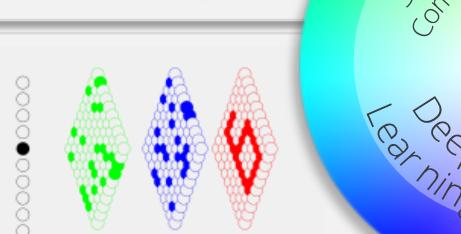
Growing impact and need

Opportunity:

*Mapping to Neuromorphic*

Example:

- Whetstone conversion of DL for spiking architectures



## *Brain-Derived AI*

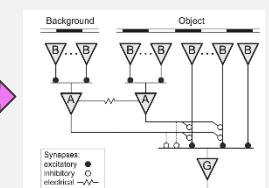
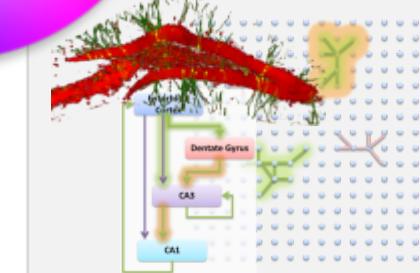
Achieve brain-like efficiency at advanced cognitive tasks, but path has proven elusive...

Opportunity:

*Develop novel algorithms that address critical DOE problems*

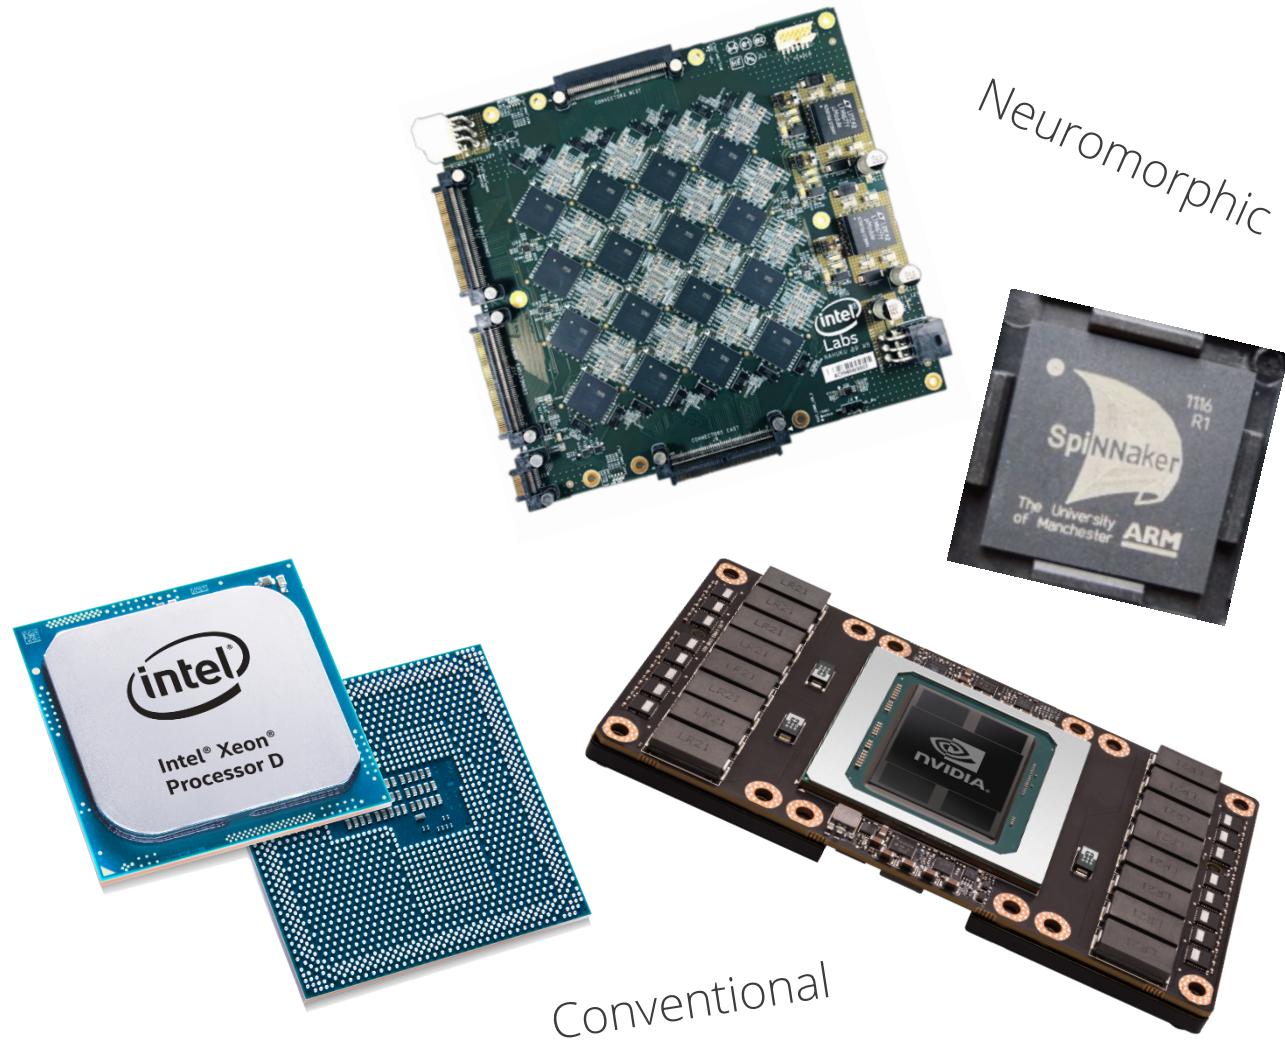
Examples:

- Dragonfly model for interception
- Hippocampus model for context-dependent learning

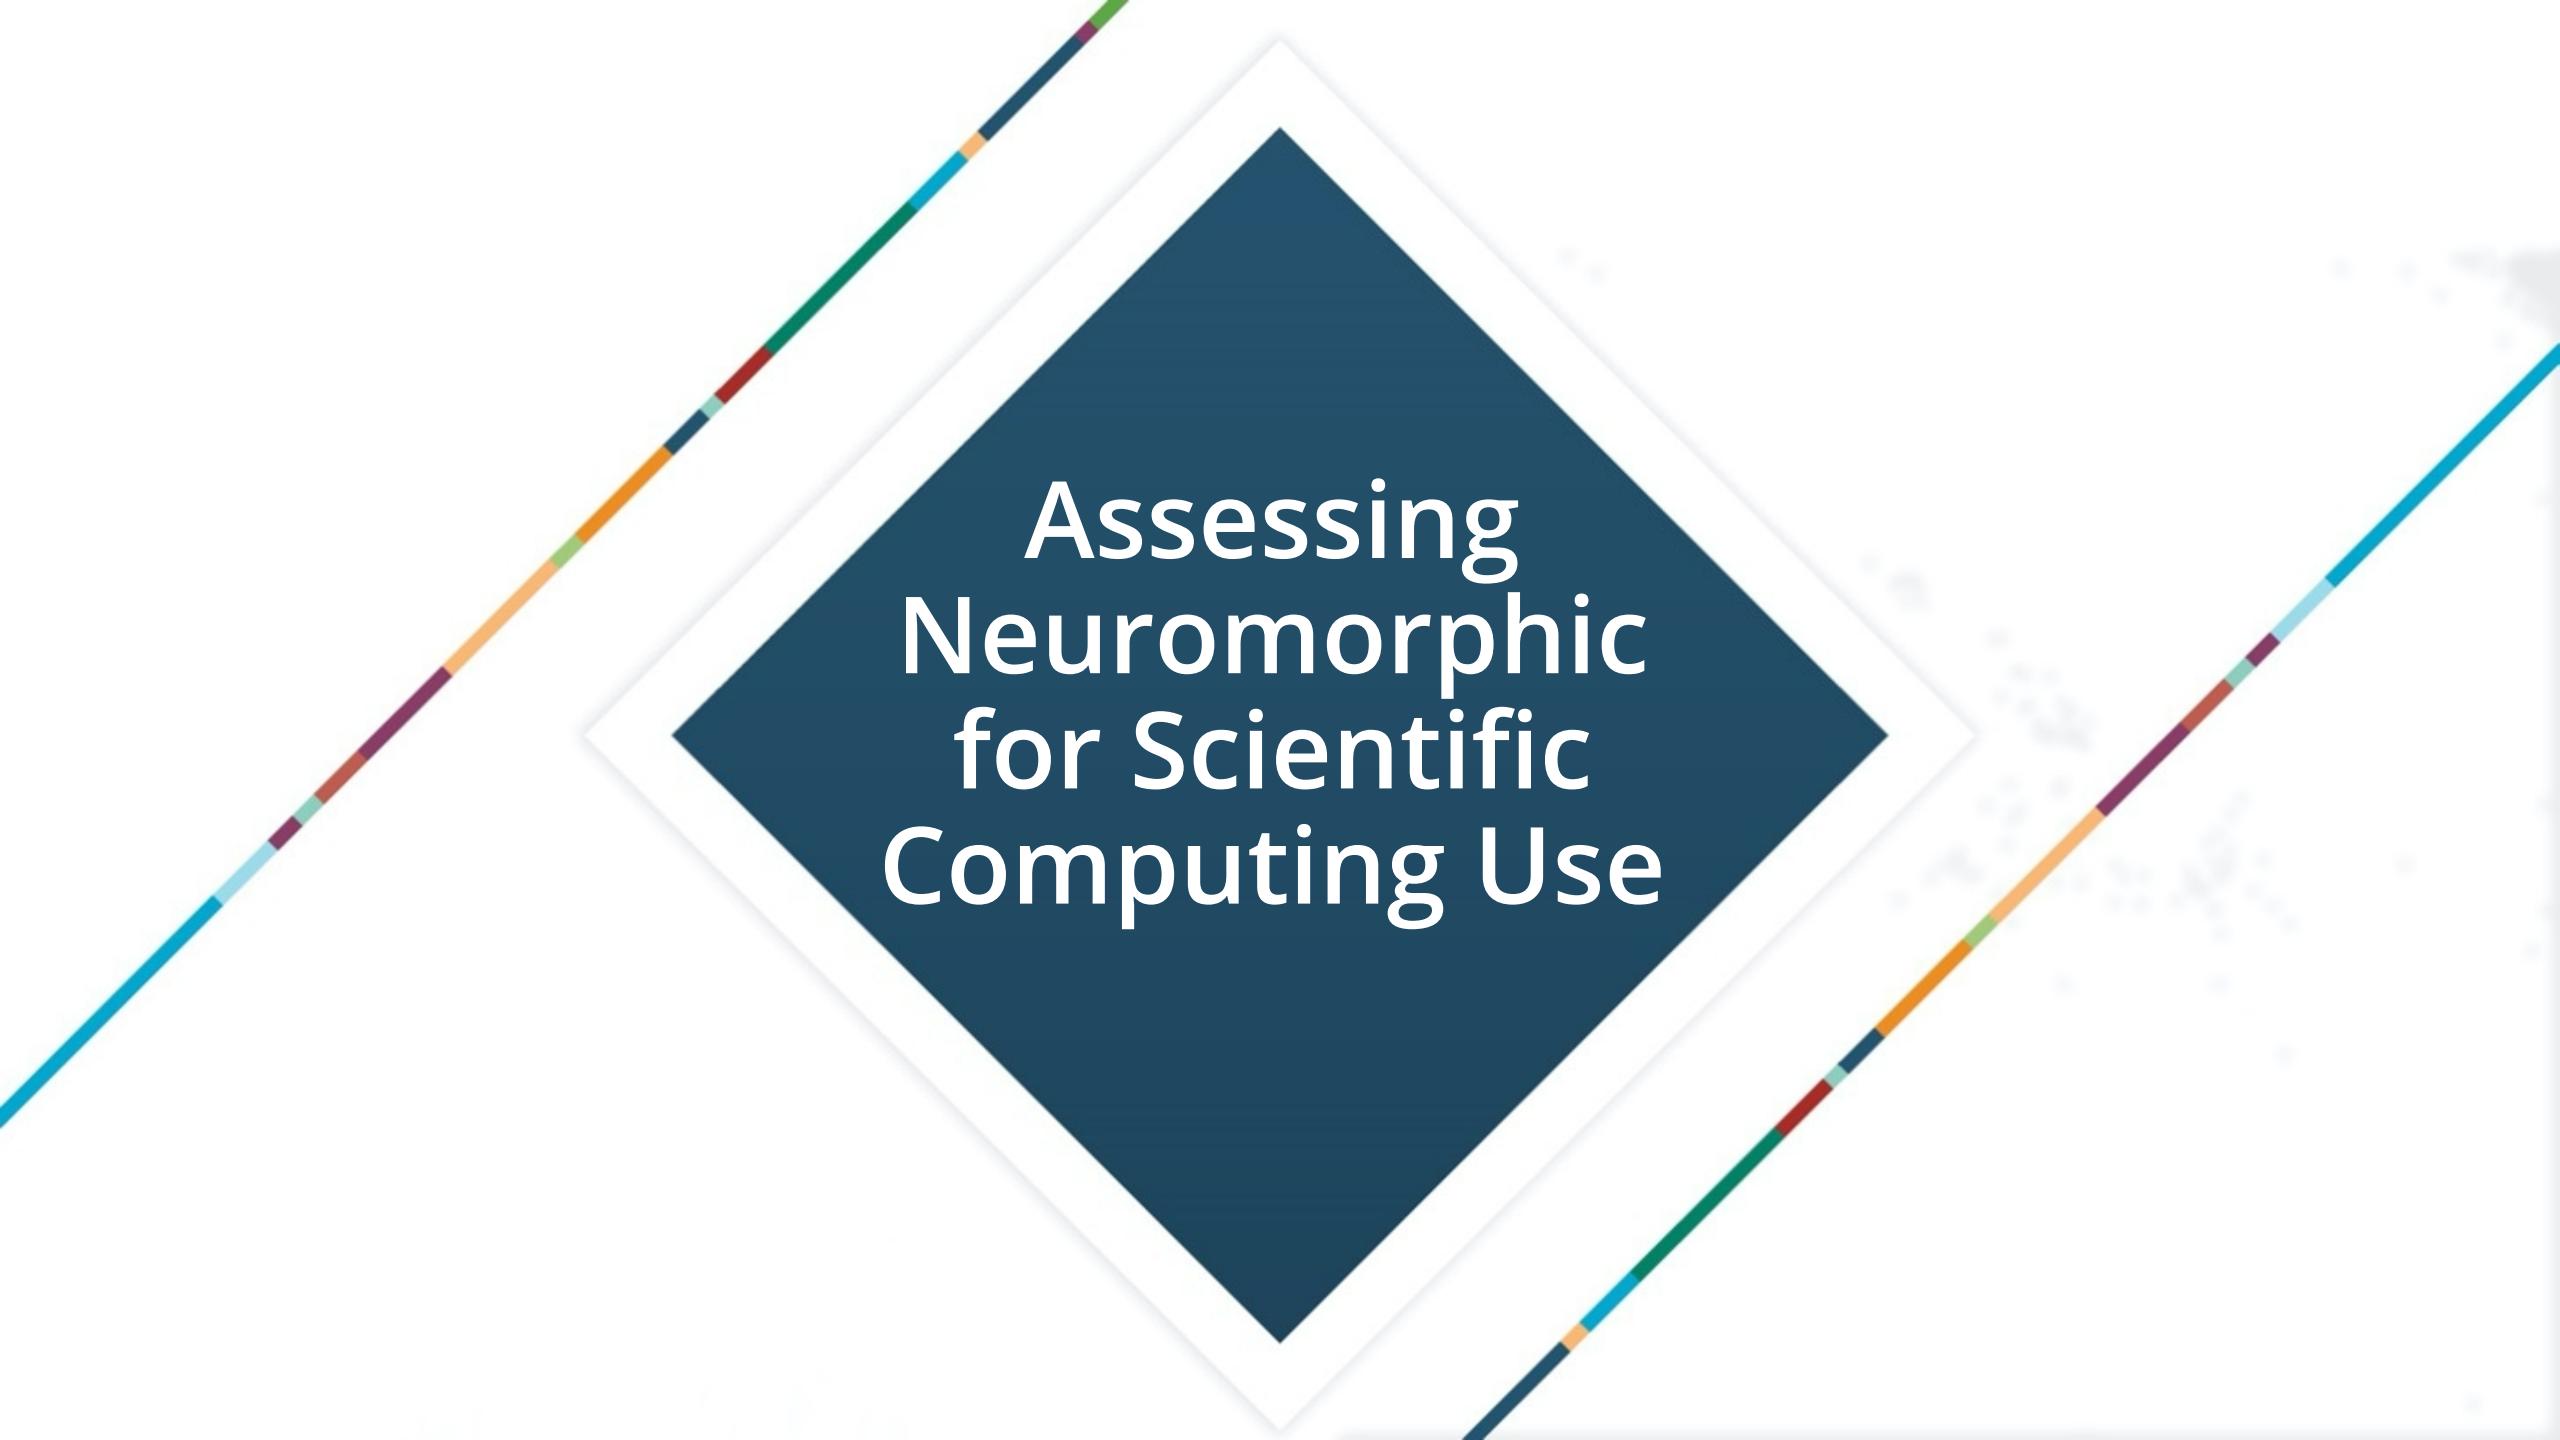


# A Heterogeneous Future

- The future of computing is likely more heterogeneous.
- Neuromorphic hardware can work side-by-side with CPUs and GPUs.
  - Scalability suggests both edge and HPC impact
- AI applications and bio-inspired algorithms will play a large part
- Energy efficiency and graph structure also potentially benefit scientific computing applications.
- We are challenged in mapping existing applications to new devices.



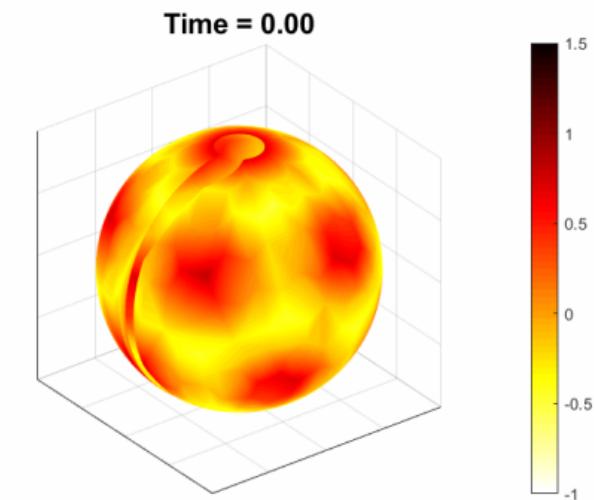
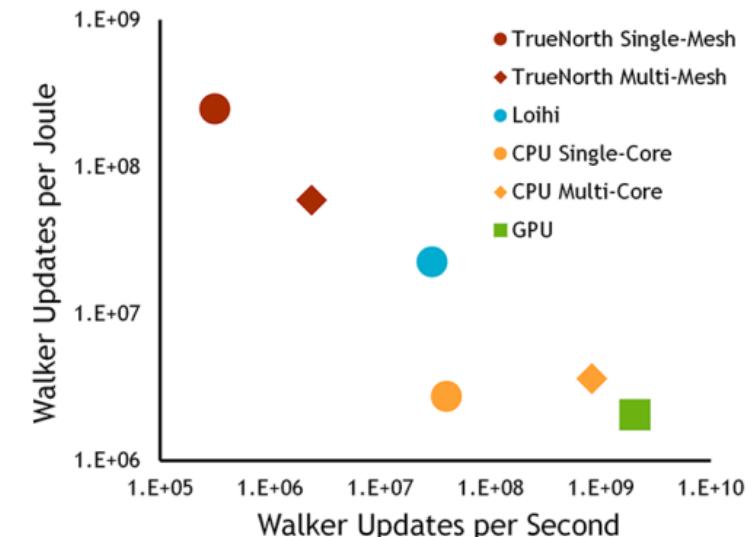
Conventional



# Assessing Neuromorphic for Scientific Computing Use

# Markov Chain Sampling on Neuromorphic

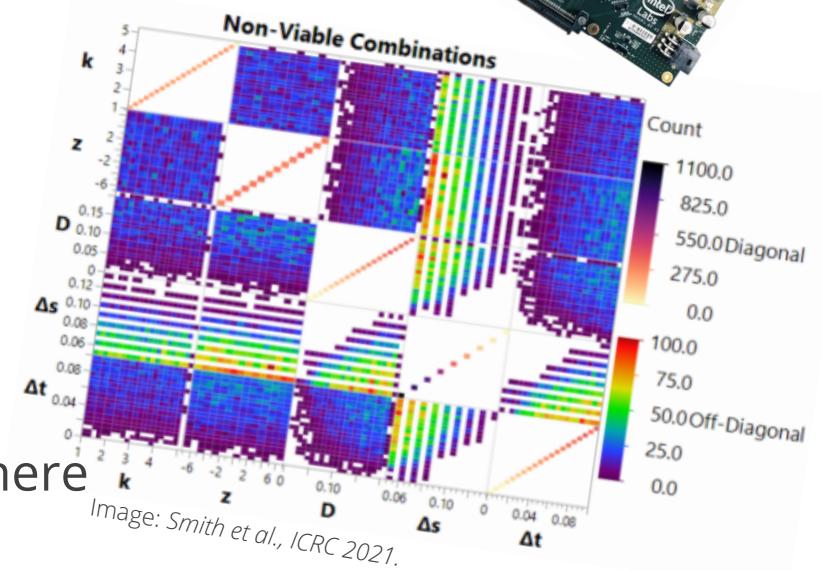
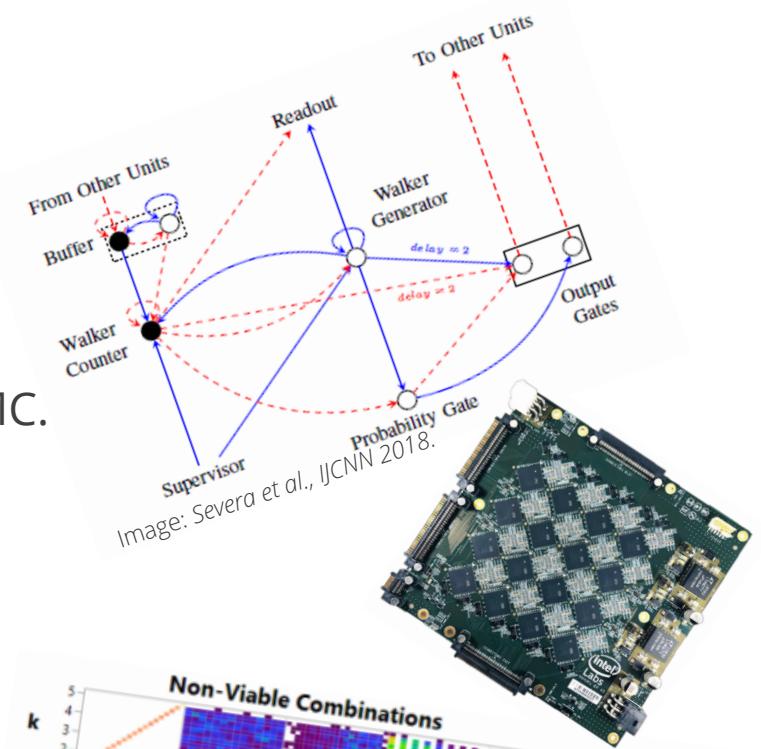
- Discrete time Markov chains (DTMCs) can be sampled efficiently on neuromorphic.
- A random walk is interpreted as a trajectory from a DTMC.
- Walkers are spikes and are routed through nodes – clusters of neurons – representing the state space of the DTMC.
- Samples obtained can be used to solve numerous PDEs, including heat and particle transport equations.





# Evaluating Sampling on Loihi

- That Loihi numerically solves PDEs well is interesting.
  - 8-bit limited PRNG.
  - Algorithm only approximates required stochastic process with a DTMC.
- **Is Loihi sampling as expected**, or are the test cases robust to stochastic process perturbations?
- **How do we assess sampling and identify where limitations may exist?**
  - How can we perform code verification and hardware validation?
- **Relative entropy** can measure how far samples are from the expected distributions and be framed as a hypothesis test.
- Further statistical analysis can explore the parameter space where the neuromorphic sampling algorithm can fail.



## Ornstein-Uhlenbeck – A Test Problem

- We assess our Loihi algorithm on the Ornstein-Uhlenbeck Equation:

$$dX(t) = -k(X(t) - z)dt + \sqrt{2D}dW(t)$$



Applications in  
Physics, Cell-bio,  
Finance, and  
Epidemiology

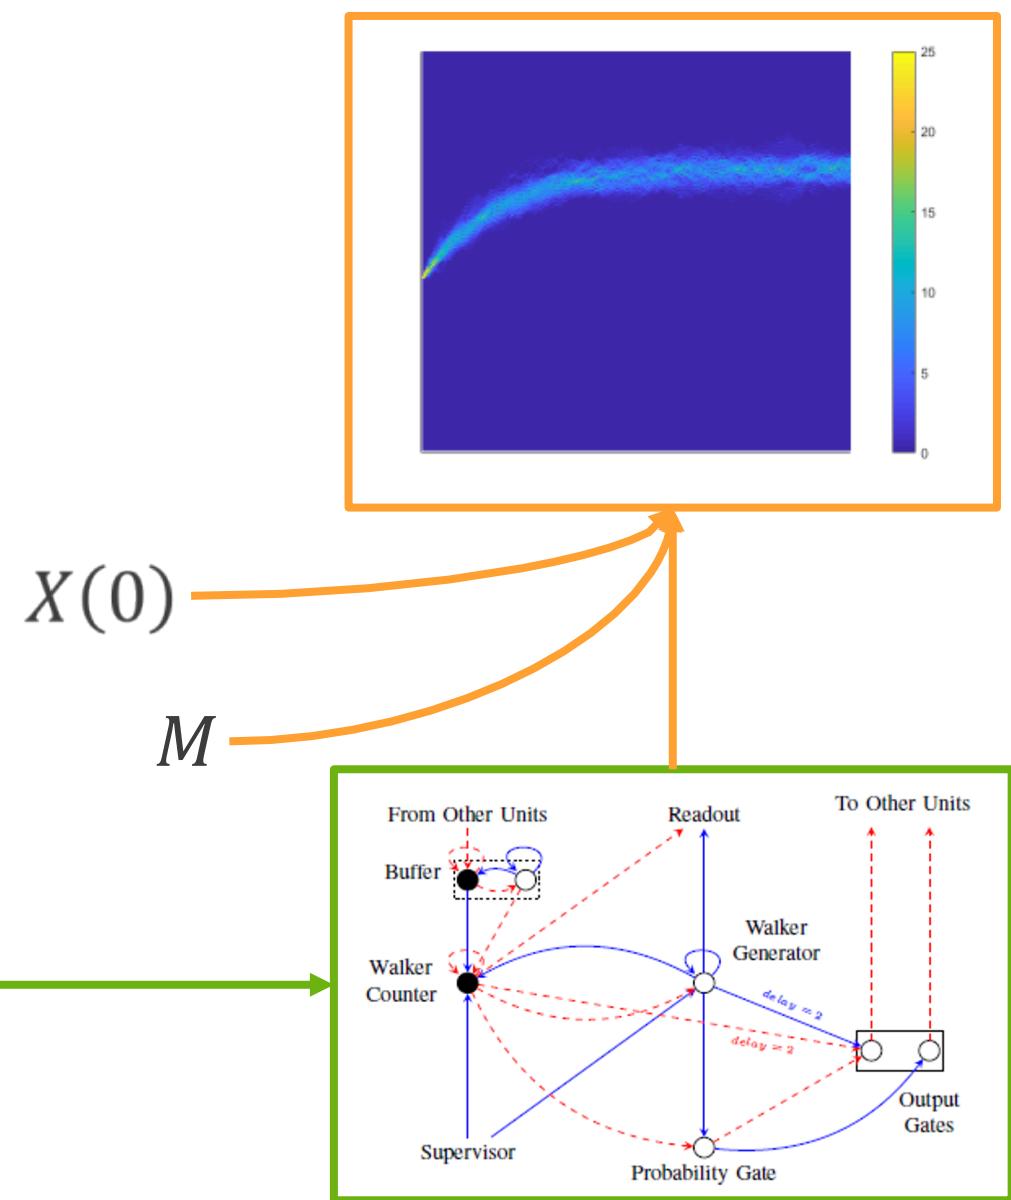
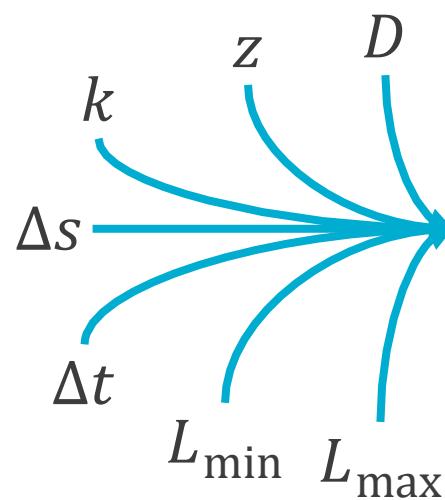
- The measure of relative entropy we select is the Kullback-Leibler Divergence:

$$\kappa(p, q) = \int p(x) \log \frac{p(x)}{q(x)} dx$$

- We let  $p$  represent the distribution of samples from Loihi and  $q$  represent the distribution of the Ornstein-Uhlenbeck process.

# Data Generation

- We generated 29.087 observations on Loihi by sampling from a DTMC given a collection of parameters.
- Observations are of  $M$  random walkers starting at location  $X(0)$ .



# Relative Entropy Results

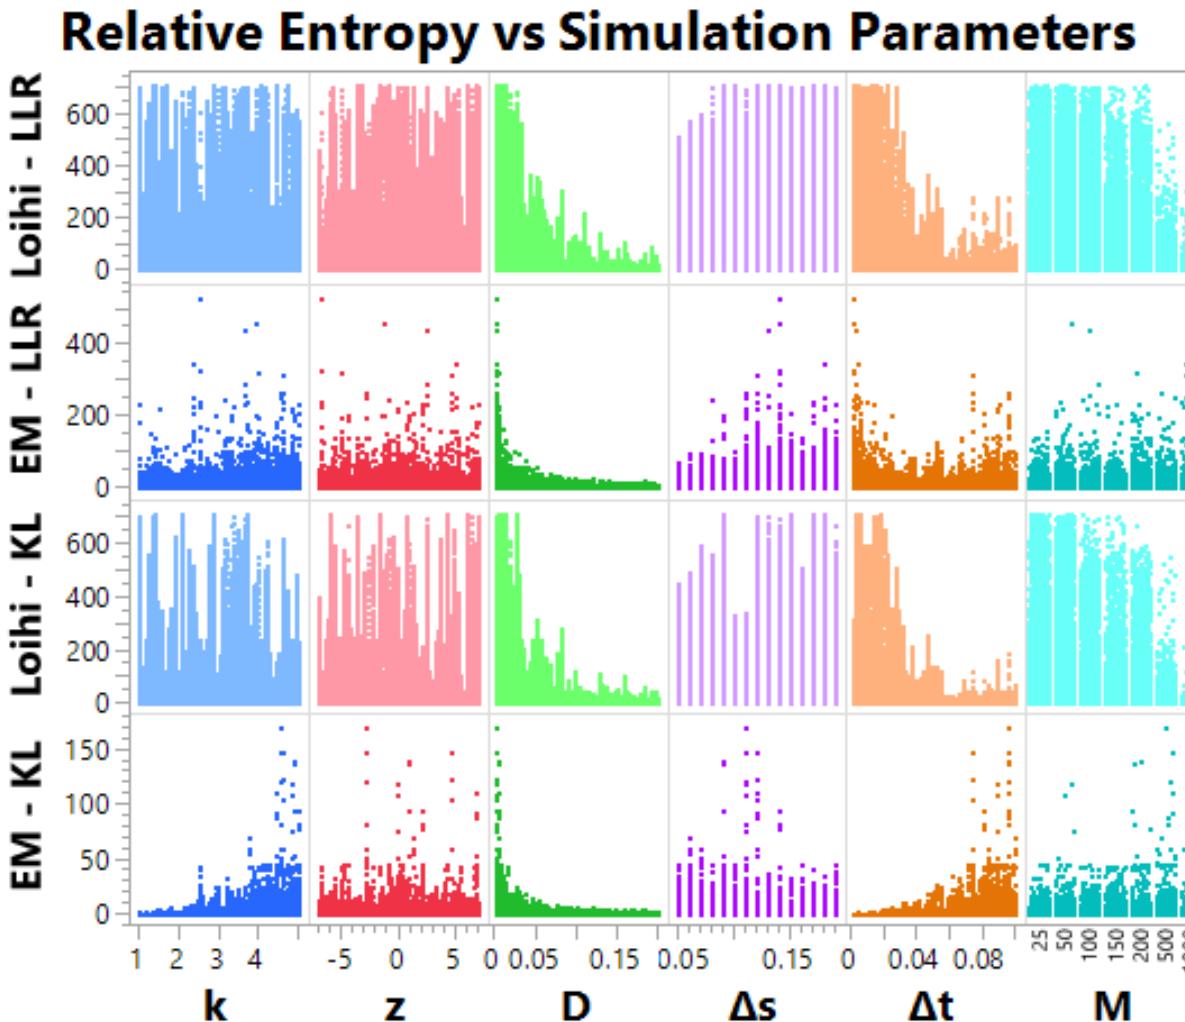


Image: Smith et al., ICRC 2021.

- In addition to the Kullback-Leibler (KL) metric, we calculated a secondary approximate metric, the Log-Likelihood Ratio (LLR).
- Additionally, we simulated trajectories in a traditional manner (Euler-Maruyama) for the exact same set of parameters.
- Each square subplot contains one dot for each time step in each of the 29,087 trajectories. This is a total of 7,715,213 total measurements per subplot.



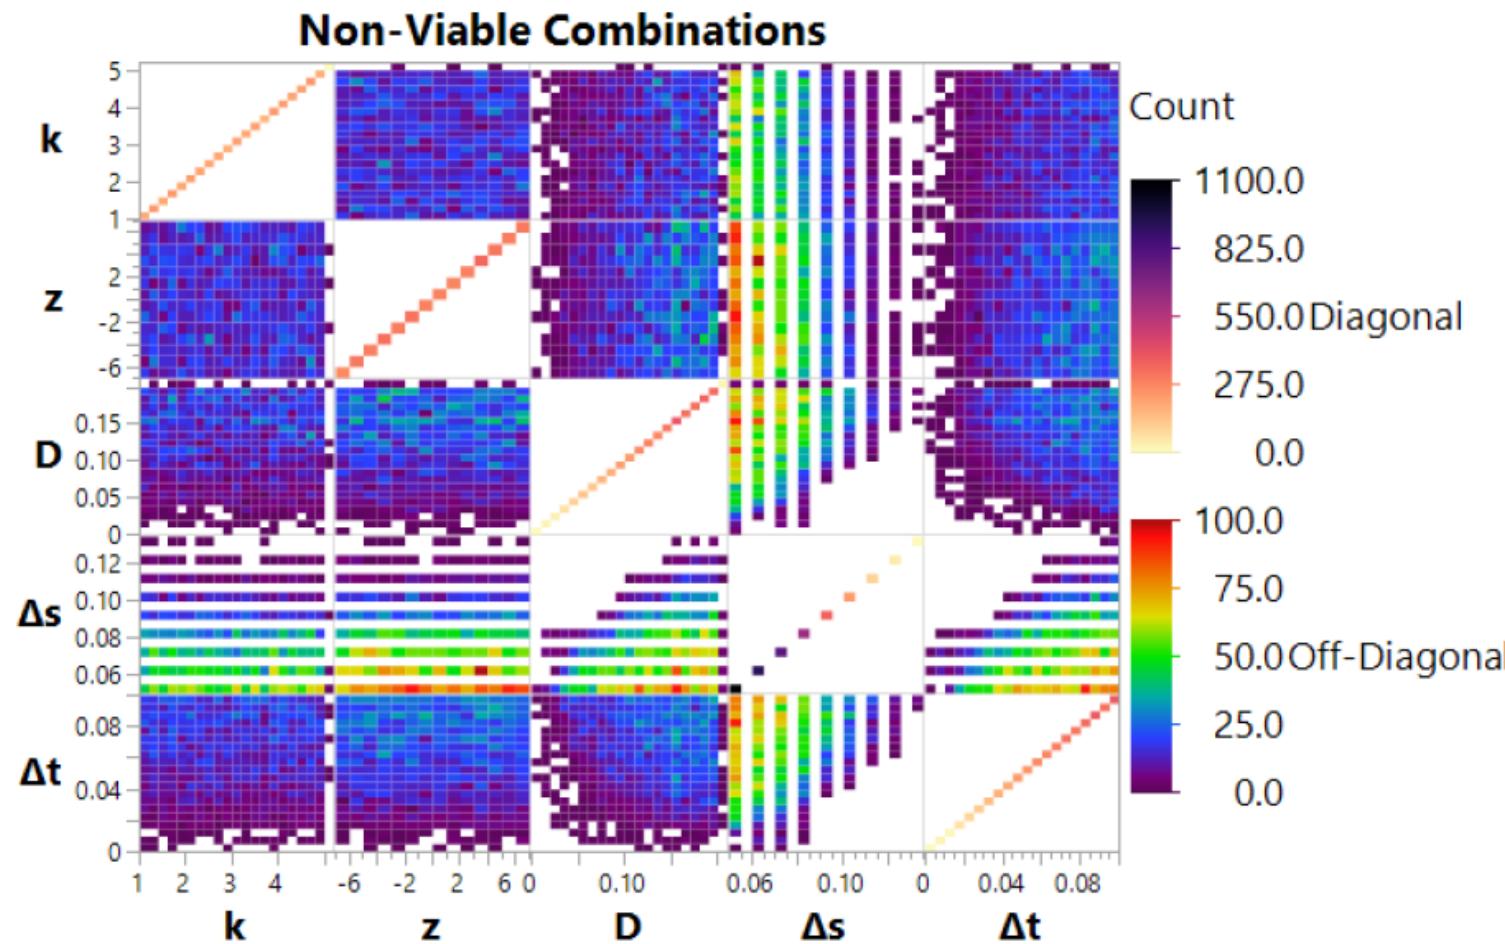
## Relative Entropy Results

- The majority of time steps observed on Loihi have a relative entropy value less than 1.
- Recalling the hypothesis testing analogy, this is highly significant. Particularly so given the large number of measurements.

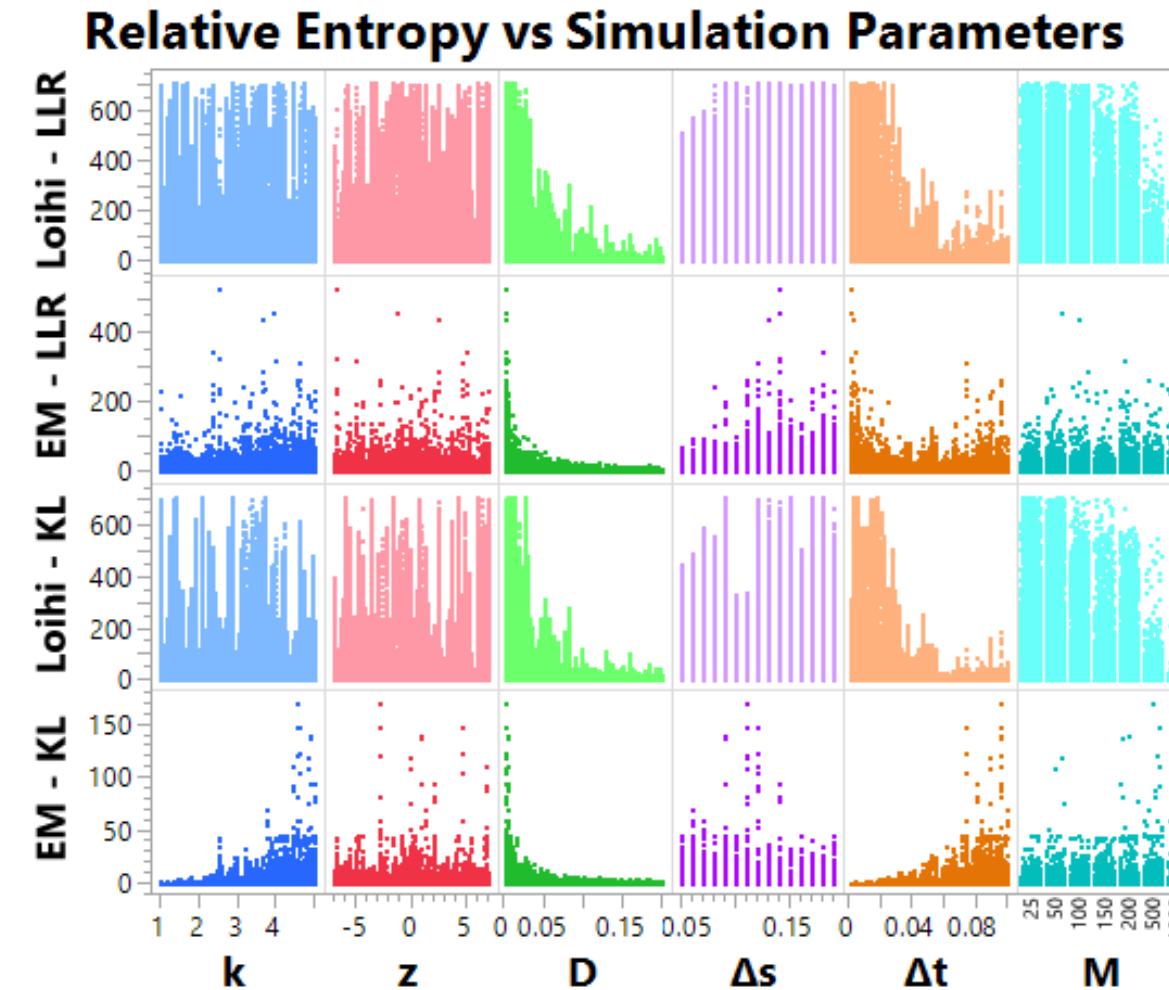
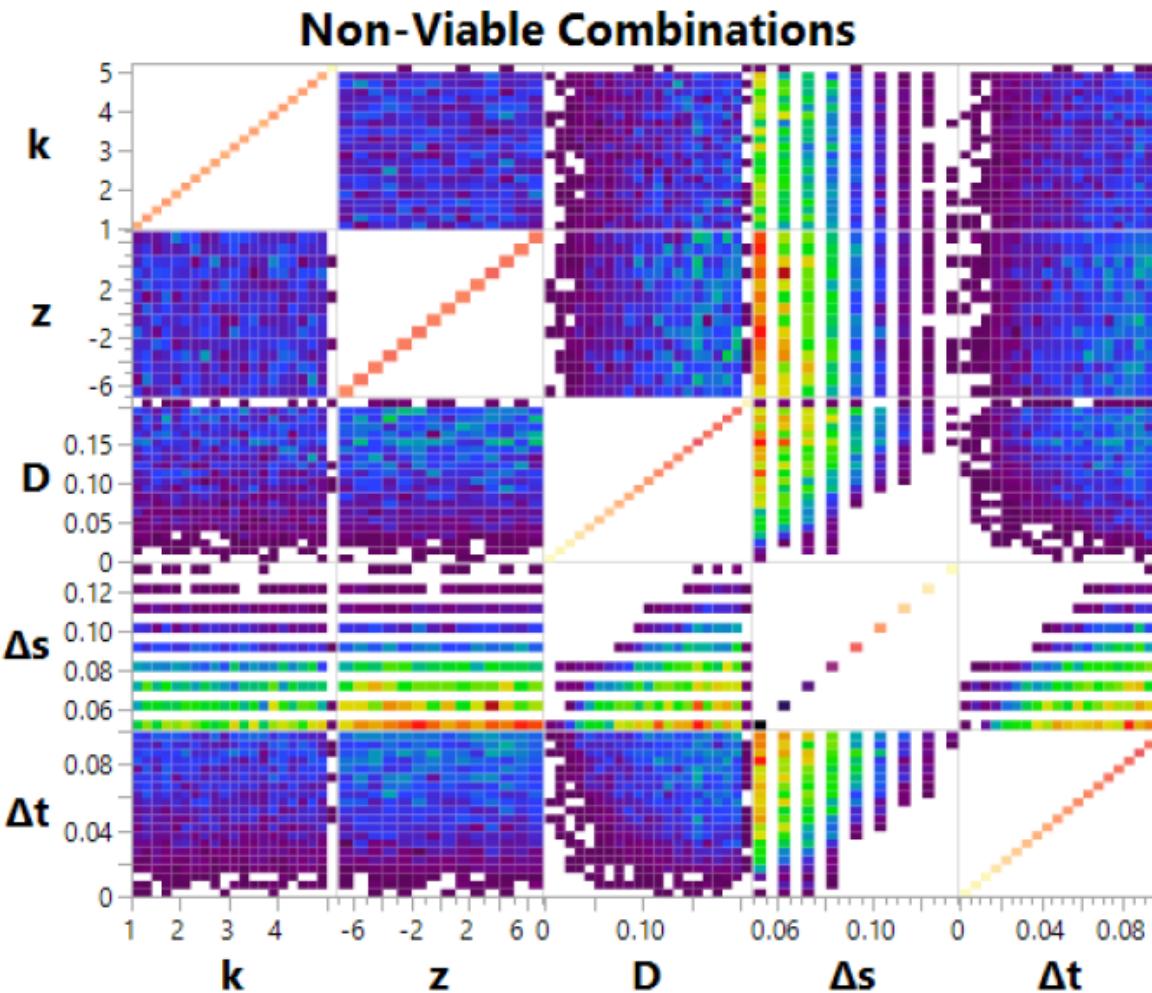
		<i>M</i>							All
		25	50	100	150	200	500	1000	
Loihi - LLR	Mean	3.0348	2.6378	1.9470	1.7794	1.4786	0.8407	0.6270	2.3859
	Std Dev	24.1788	21.8369	18.8534	15.9938	14.4001	6.9788	5.7801	20.6868
	90% Quantile	1.0880	0.8682	0.7646	0.8320	0.6899	0.7300	0.7626	0.9200
EM - LLR	Mean	0.3517	0.2719	0.2097	0.2017	0.2128	0.2535	0.3536	0.2790
	Std Dev	0.7112	1.0117	1.1081	0.9558	1.5374	1.8285	3.4872	1.1118
	90% Quantile	0.5268	0.3615	0.2853	0.2810	0.2669	0.3348	0.4801	0.4407
Loihi - KL	Mean	1.8094	1.6362	1.1634	1.0352	0.8980	0.4350	0.2031	1.4318
	Std Dev	15.4816	15.8599	11.9741	10.1031	10.5996	5.0513	1.9176	13.8093
	90% Quantile	0.5350	0.3383	0.2463	0.2503	0.2118	0.2387	0.2243	0.4264
EM - KL	Mean	0.1917	0.1105	0.0665	0.0529	0.0477	0.0545	0.0763	0.1187
	Std Dev	0.1796	0.2248	0.2020	0.2809	0.4522	0.8234	0.8567	0.3086
	90% Quantile	0.3588	0.1996	0.1118	0.0805	0.0648	0.0442	0.0837	0.2593

# Viable Sampling Space

- We recovered 4229 combinations of parameters that were unable to be sampled due to exhaustion of probabilistic resources.



# A Trade-Off in Simulation



# Where do hardware limitations end and algorithm limitations begin?

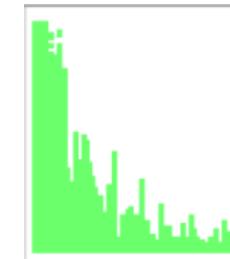
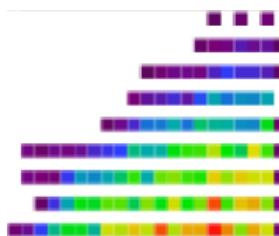
## Hardware

Limited PRNG availability and precision.

$\frac{1}{256}$



Viable parameter space limitations and accuracy implications.

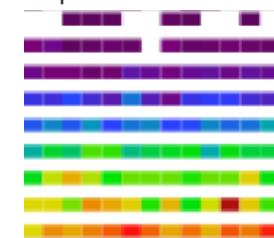


## Algorithm

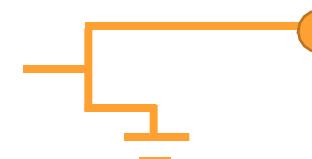
- More resource-intensive probability circuits.



- Further limited parameter space.



- Clever efficient circuits?

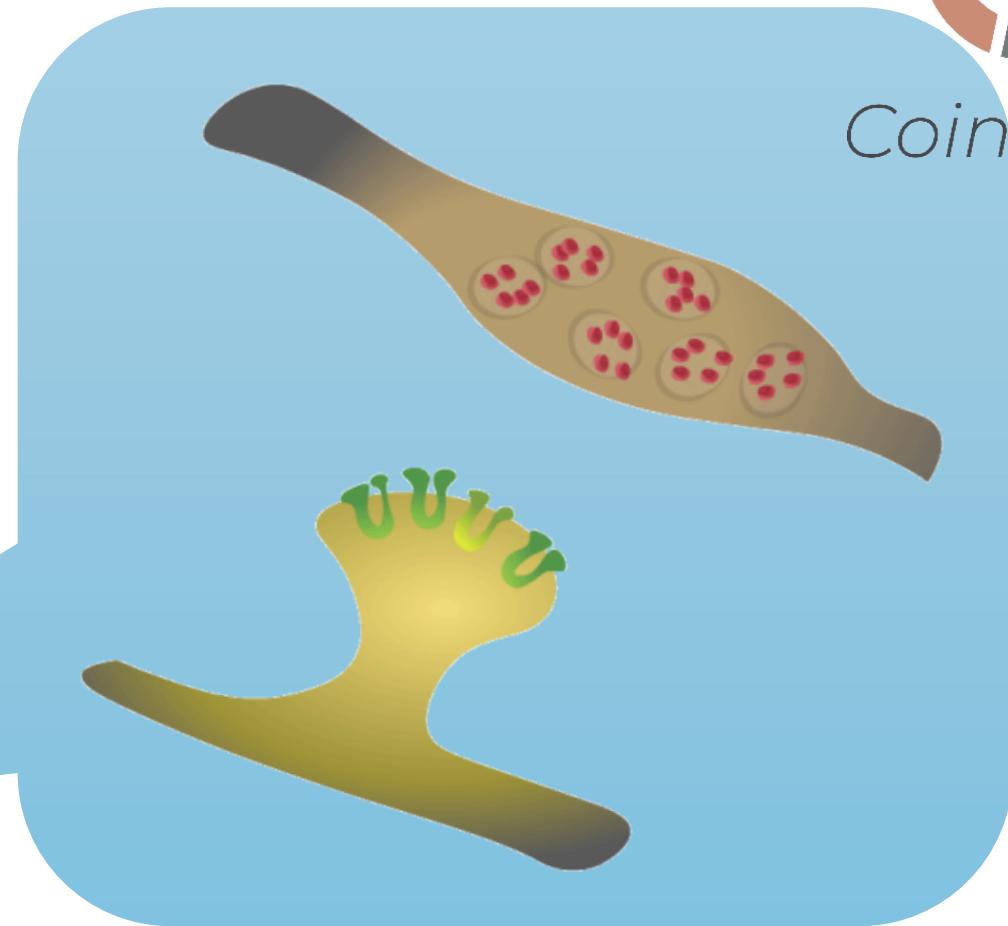
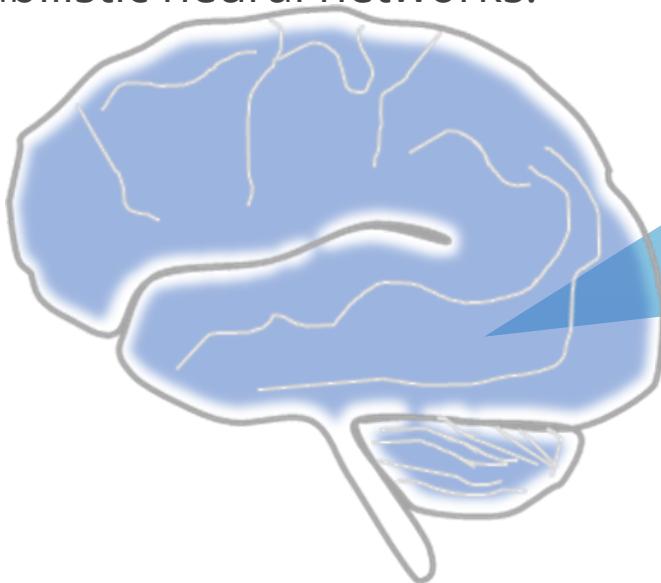




# Sampling Validation as a Methodology

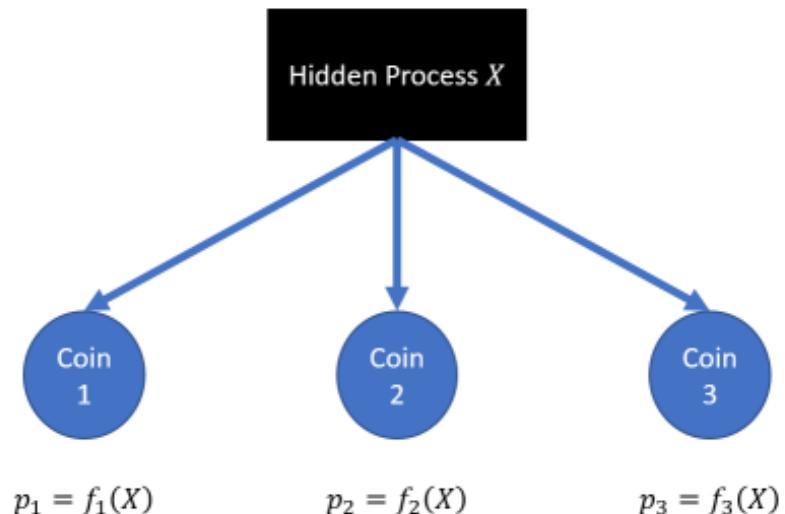
# Assessing New Devices, From the Ground Up

- The goals of the DOE Office of Science COINFLIPS project – Co-Designed Improved Neural Foundations Leveraging Inherent Physics Stochasticity:
  - Make stochasticity ubiquitous and useful;
  - Directly sample from distributions of interest;
  - Allow more efficient sampling of probabilistic neural networks.



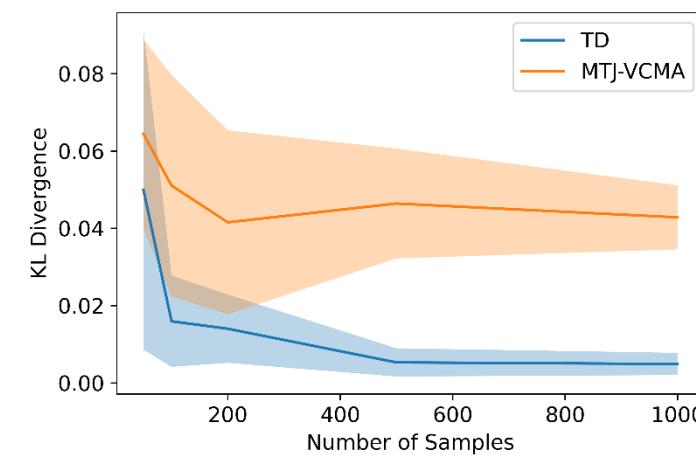
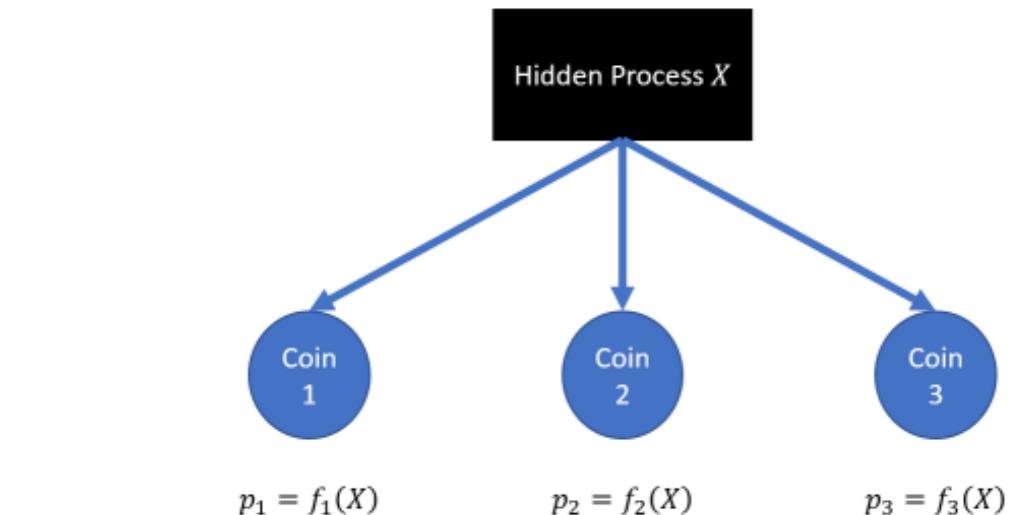
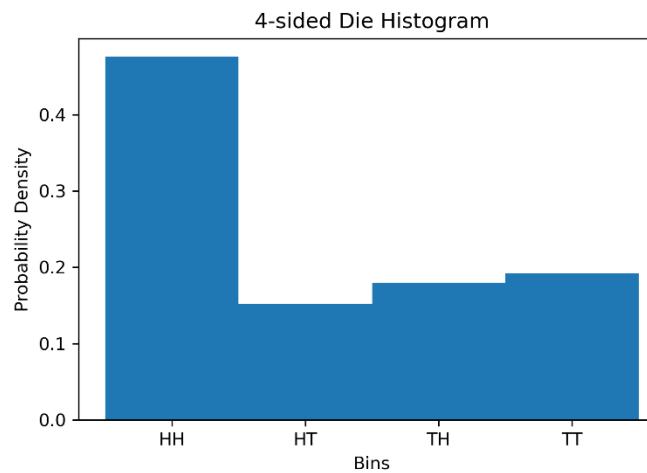
# Assessing New Devices, From the Ground Up

- The COINFLIPS team are examining two ways of generating random coin flips, one a Magnetic Tunnel Junction, and the other a Tunnel Diode.
- The devices can be combined to obtain distributions of interest. For instance, a weighted die that rolls 1 half of the time, and the numbers 2, 3, and 4 equally likely the rest of the time.
- This can be represented as the outcome of two coins with probabilities that change through a hidden process.



# Assessing New Devices, From the Ground Up

- Assessment of samples and validation of hardware can again be performed through a hypothesis test.
- This time, the alternative hypothesis is that the device circuit follows the expected/desired distribution, and the null hypothesis is that it behaves any other way.





## Concluding Thoughts, Thanks!

- Mapping applications to new hardware requires critical thoughts on both code verification and hardware validation.
- New devices built from the ground up require equal consideration for how they will meet the needs of modern applications.
- Hypothesis testing and relative entropy, as a framework, provide one method of statistically rigorous assessment.