Mosaics

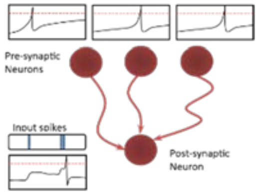

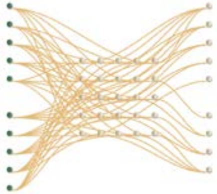

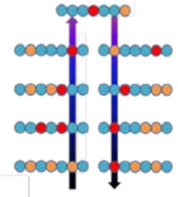

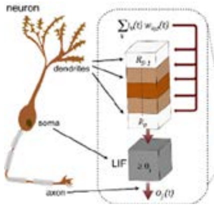

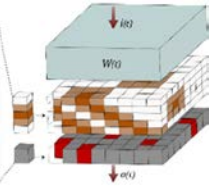

PRESENTED BY

Brad Aimone

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.

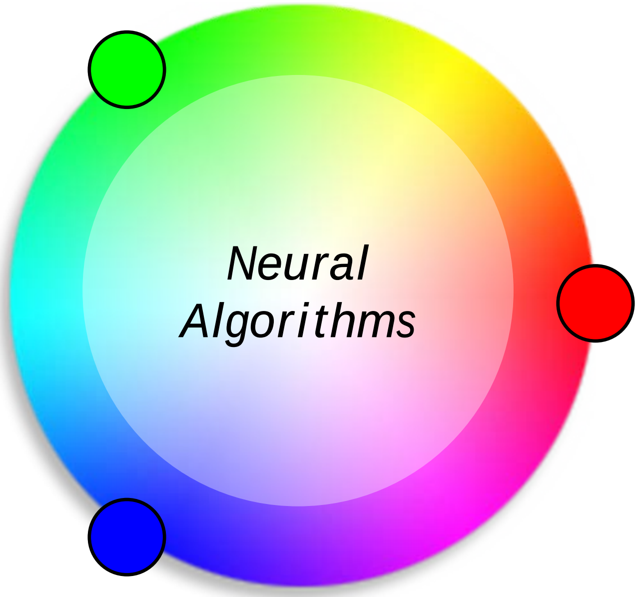

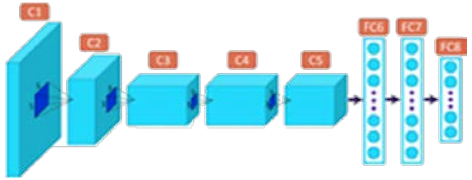

Artificial neural networks

- Generic layers of non-linear nodes
- SGD optimization of weights
- Powerful machine learning capabilities through learning sequential non-linear mappings and function approximation

Spiking neural algorithms

- Hand-crafted circuits of spiking neurons
- Model of parallel computation
- Energy efficiency through event-driven communication and high fan-in logic

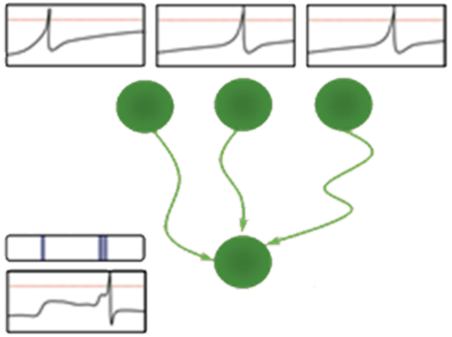

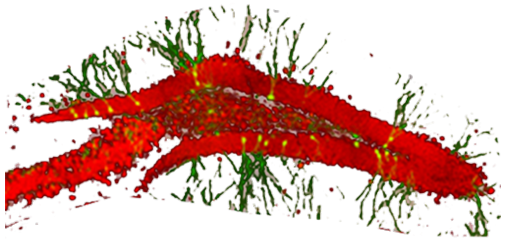

Neuroscience-constrained algorithms

- Circuit architecture based on local and regional neural connectivity
- Computation incorporates broad range of neural plasticity and dynamics
- *Generally still unexplored from algorithms perspective*

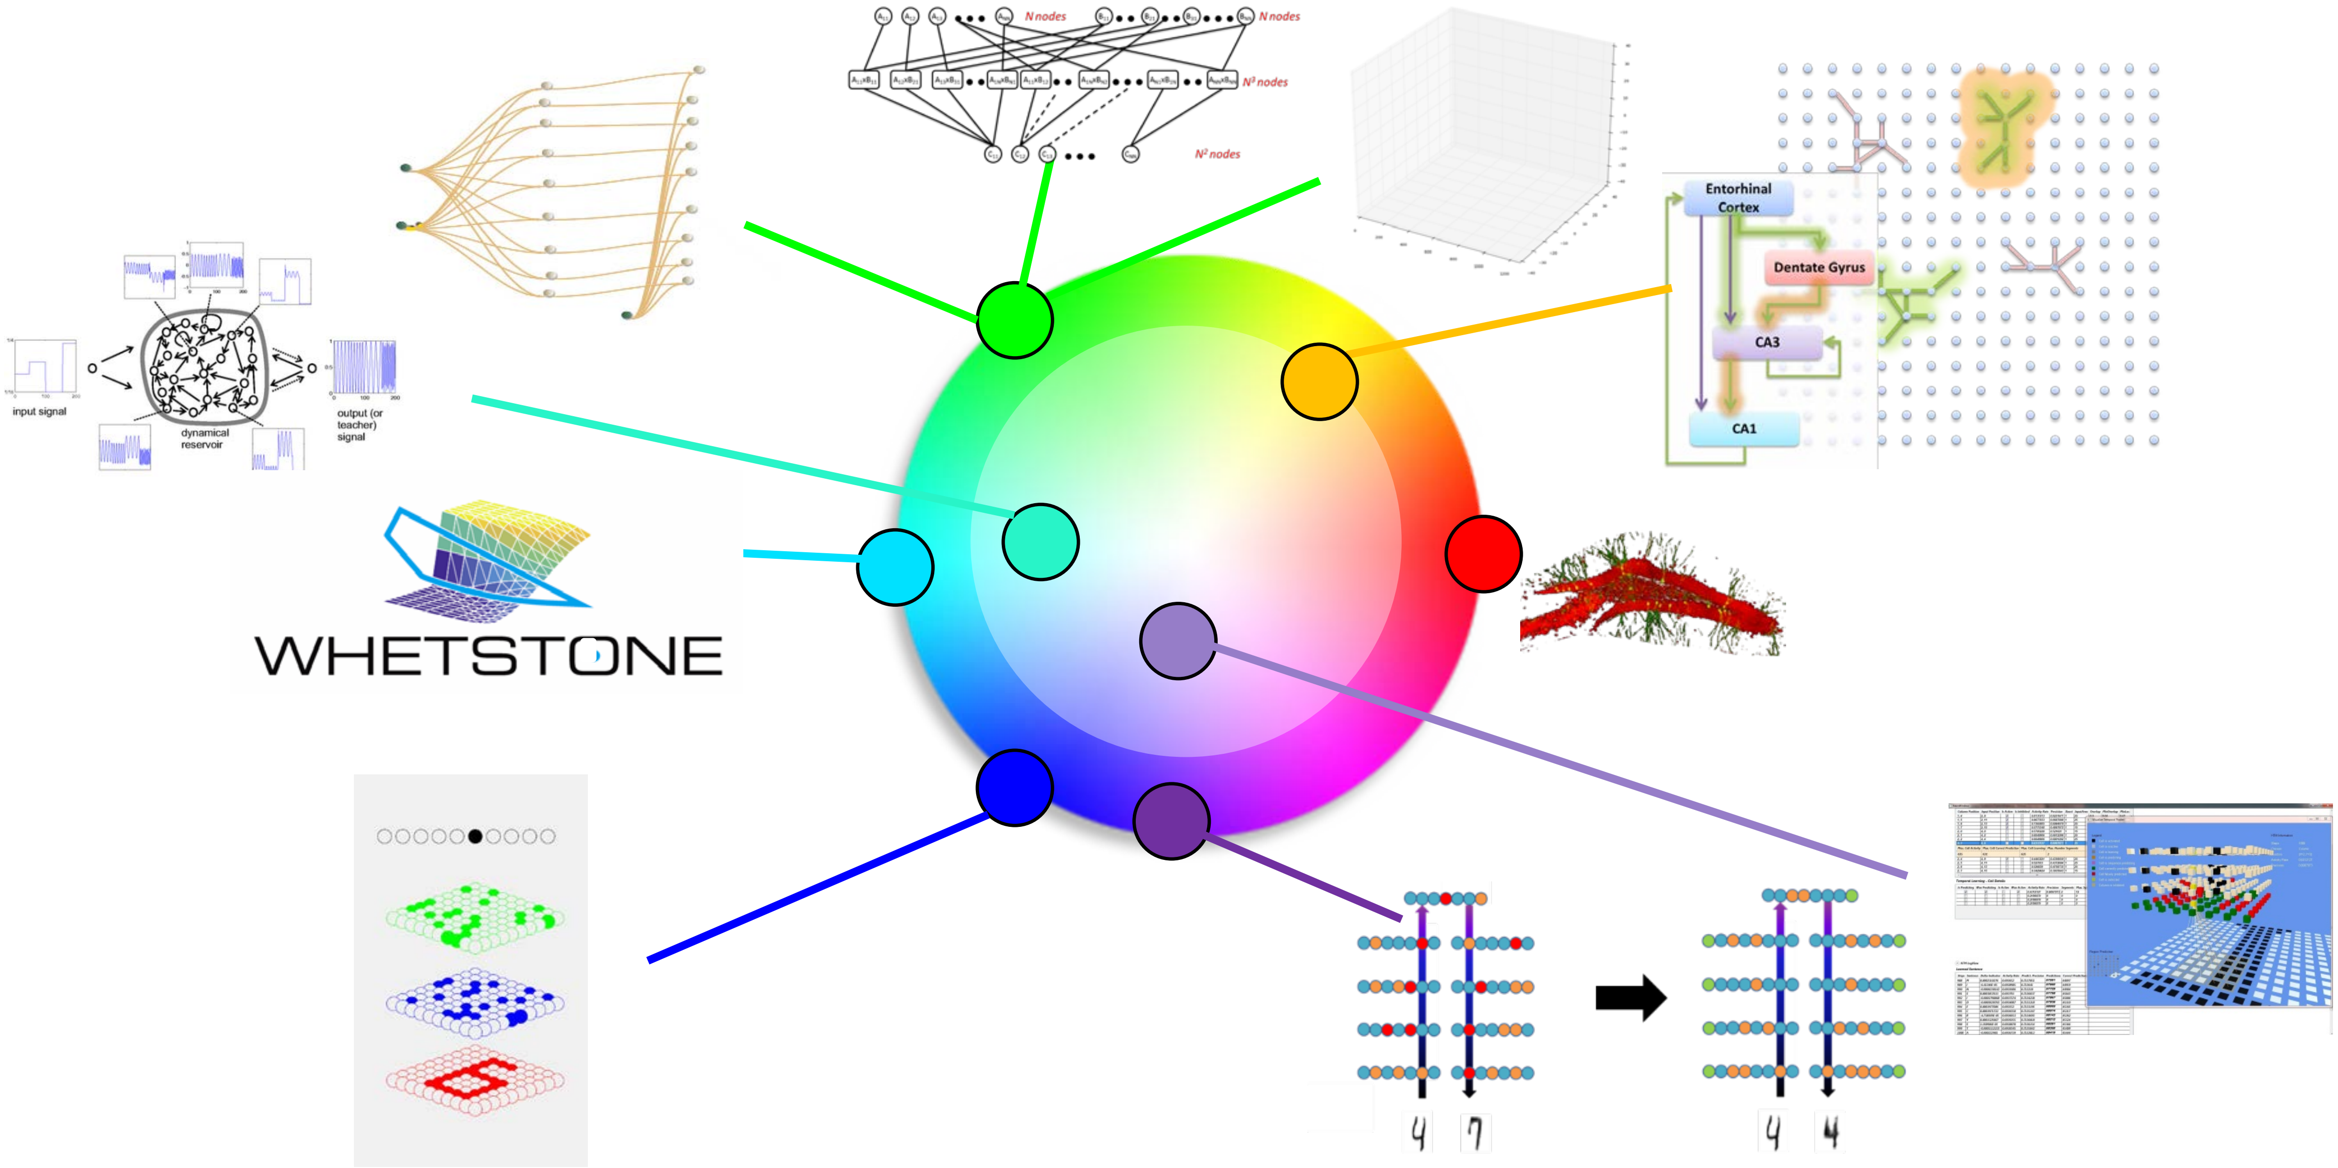
N nodes
N nodes
N³ nodes
N² nodes
Entorhinal Cortex
Dentate Gyrus
CA3
CA1
input signal
dynamical reservoir
output (or teacher) signal
WHETSTONE

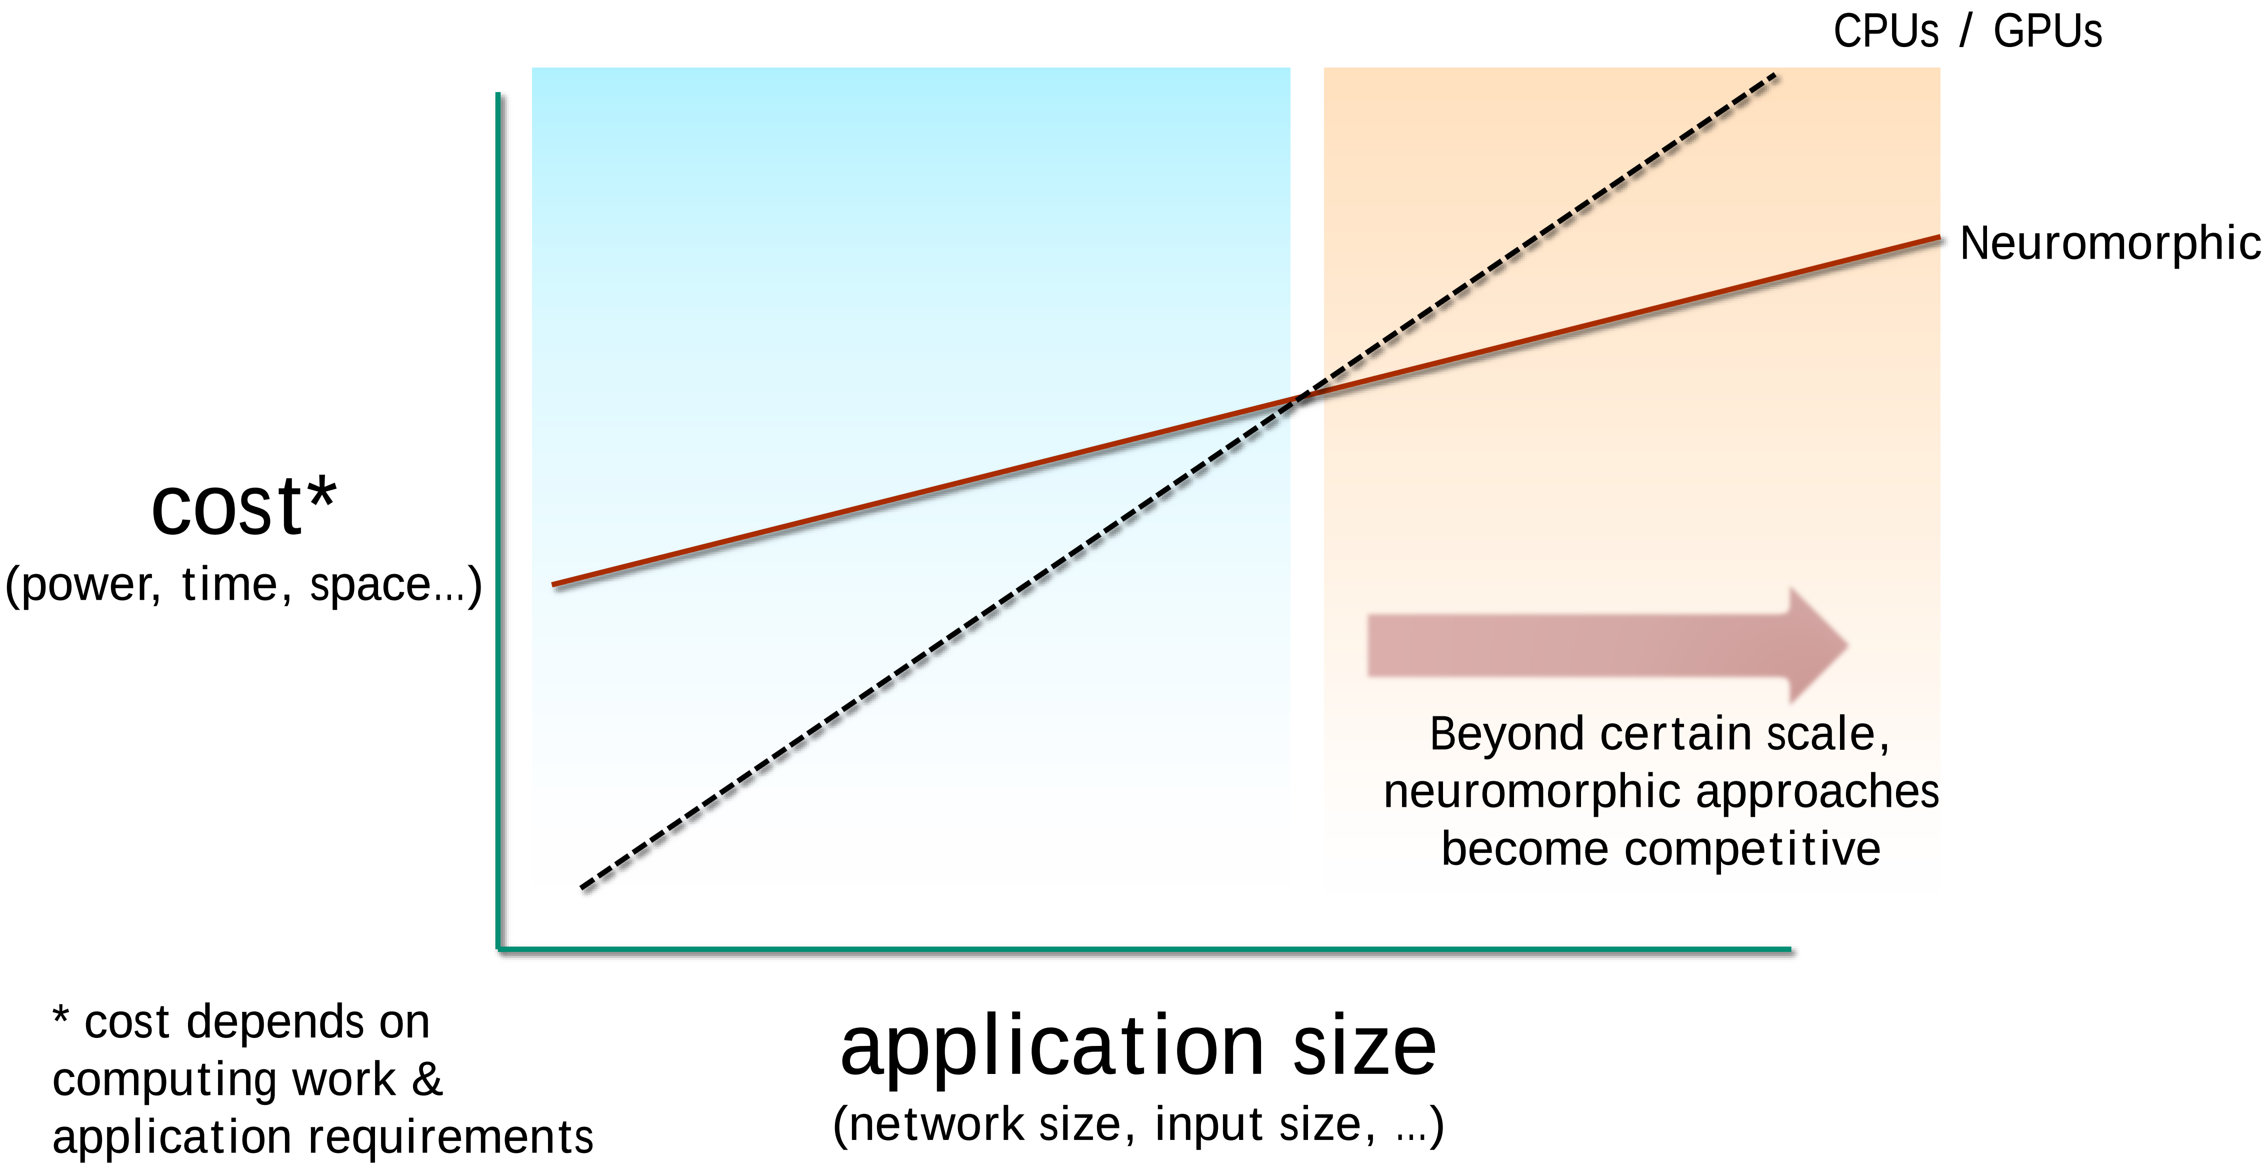
CPUs / GPUs
Neuromorphic
cost*
(power, time, space…)
Beyond certain scale,
neuromorphic approaches
become competitive
application size
(network size, input size, …)
* cost depends on
computing work &
application requirements

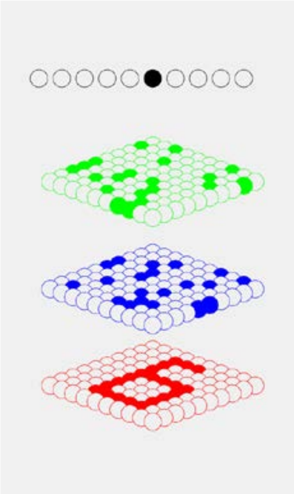

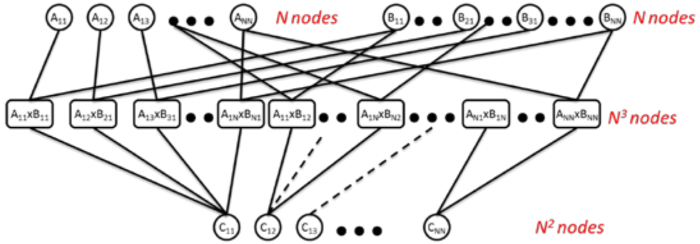

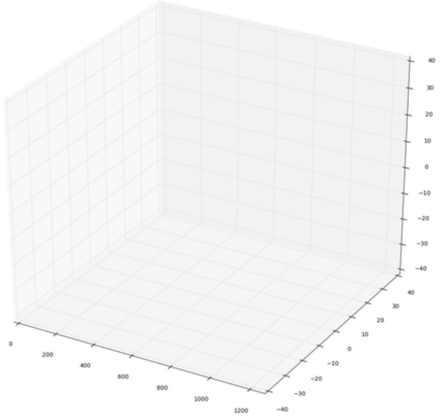

Neural Networks	Matrix Multiplication	Partial Differential Equations
• Better-suited for neuromorphic hardware than many other machine learning techniques • ANNs only became broadly world class when they reached substantially large sizes	• Neural algorithms can improve implementation of Strassen-techniques • Strassen techniques only make sense for large matrix multiplications	• Neural algorithms can efficiently implement Monte Carlo solutions for solving diffusion-based PDEs • Monte Carlo methods make most sense for high-dimensional PDEs

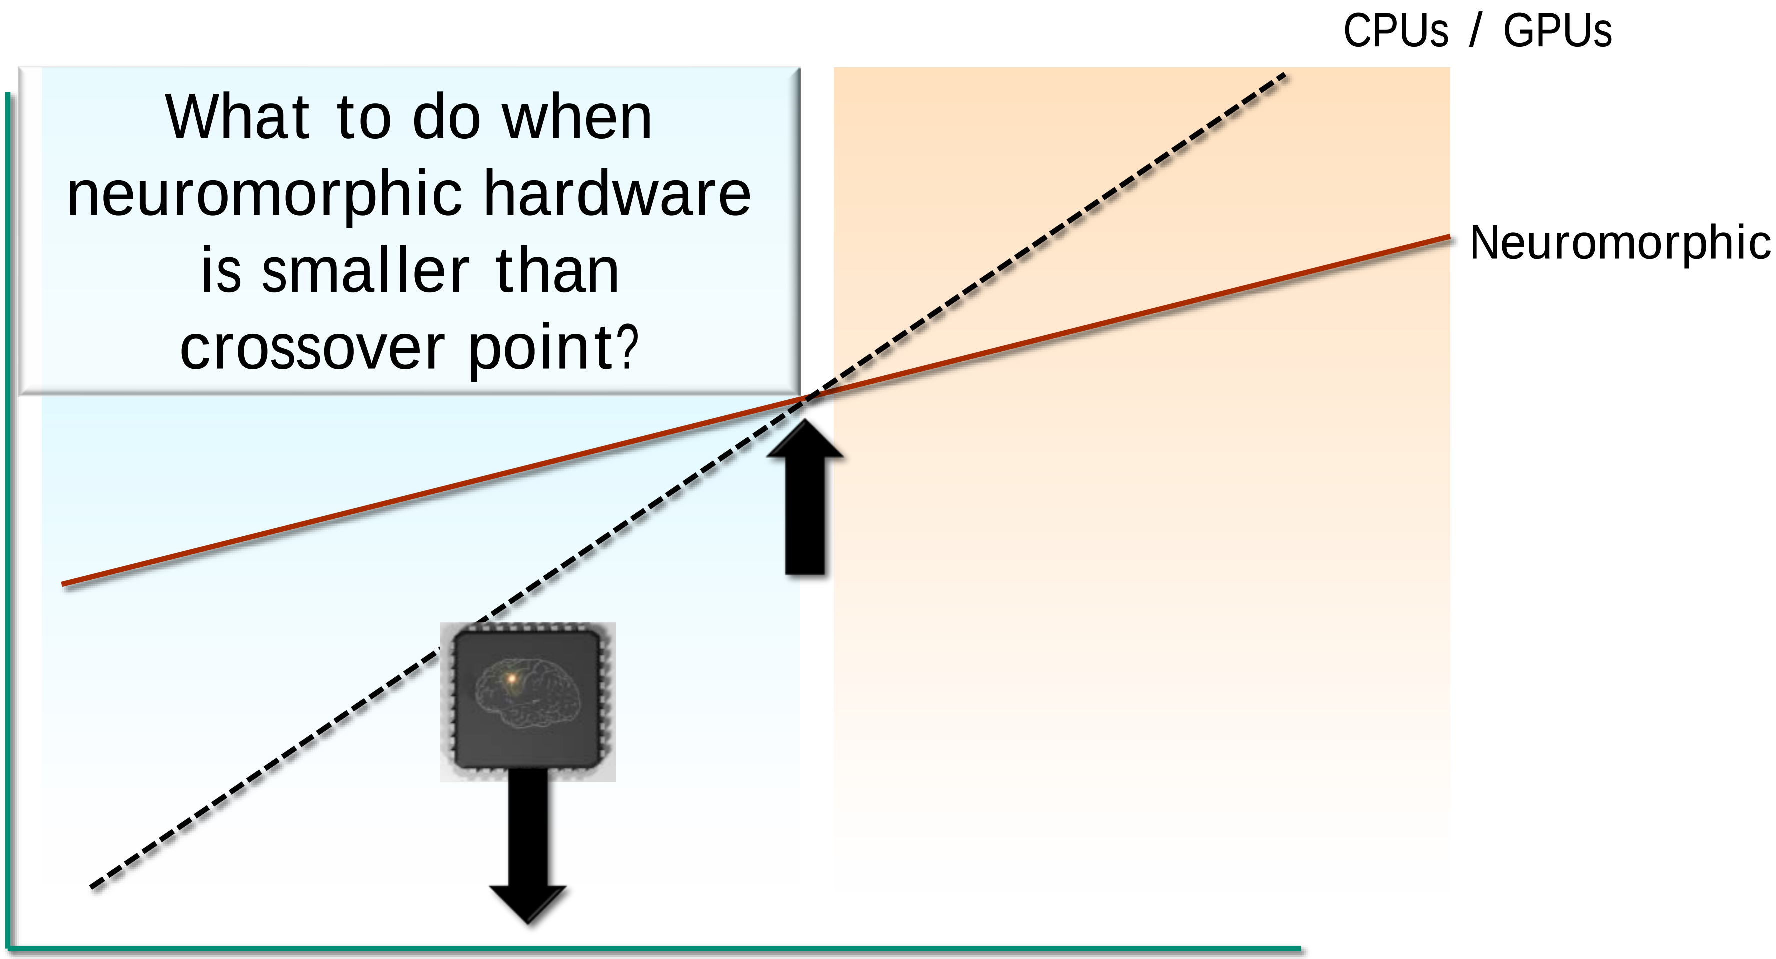

cost*

(power, time, space...)

application size

(network size, input size, ...)

* cost depends on computing work & application requirements

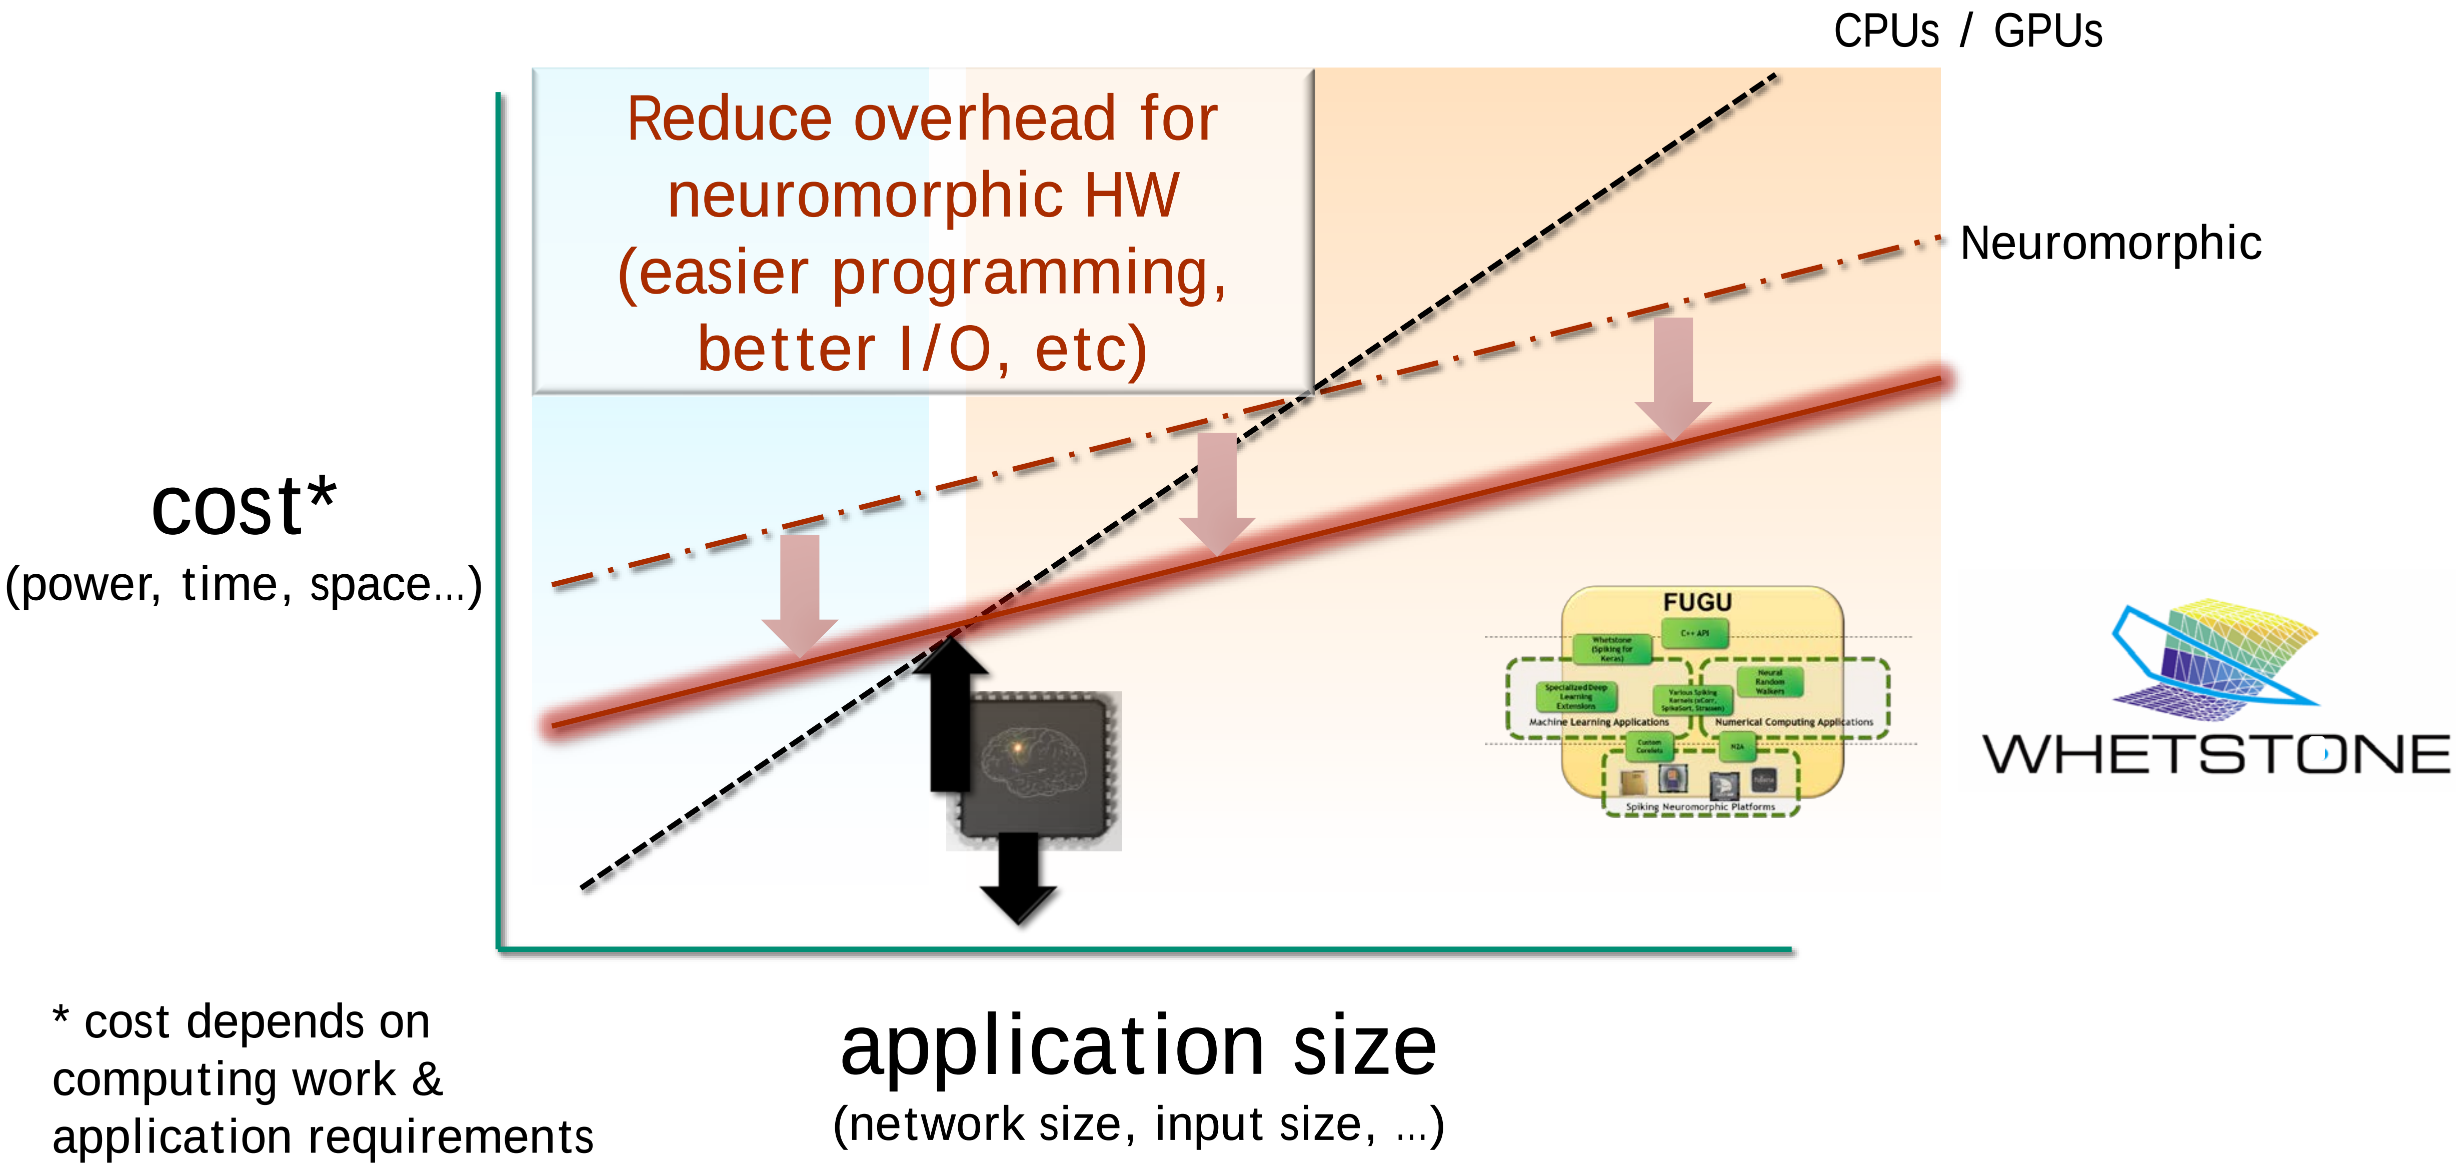

* cost depends on computing work & application requirements

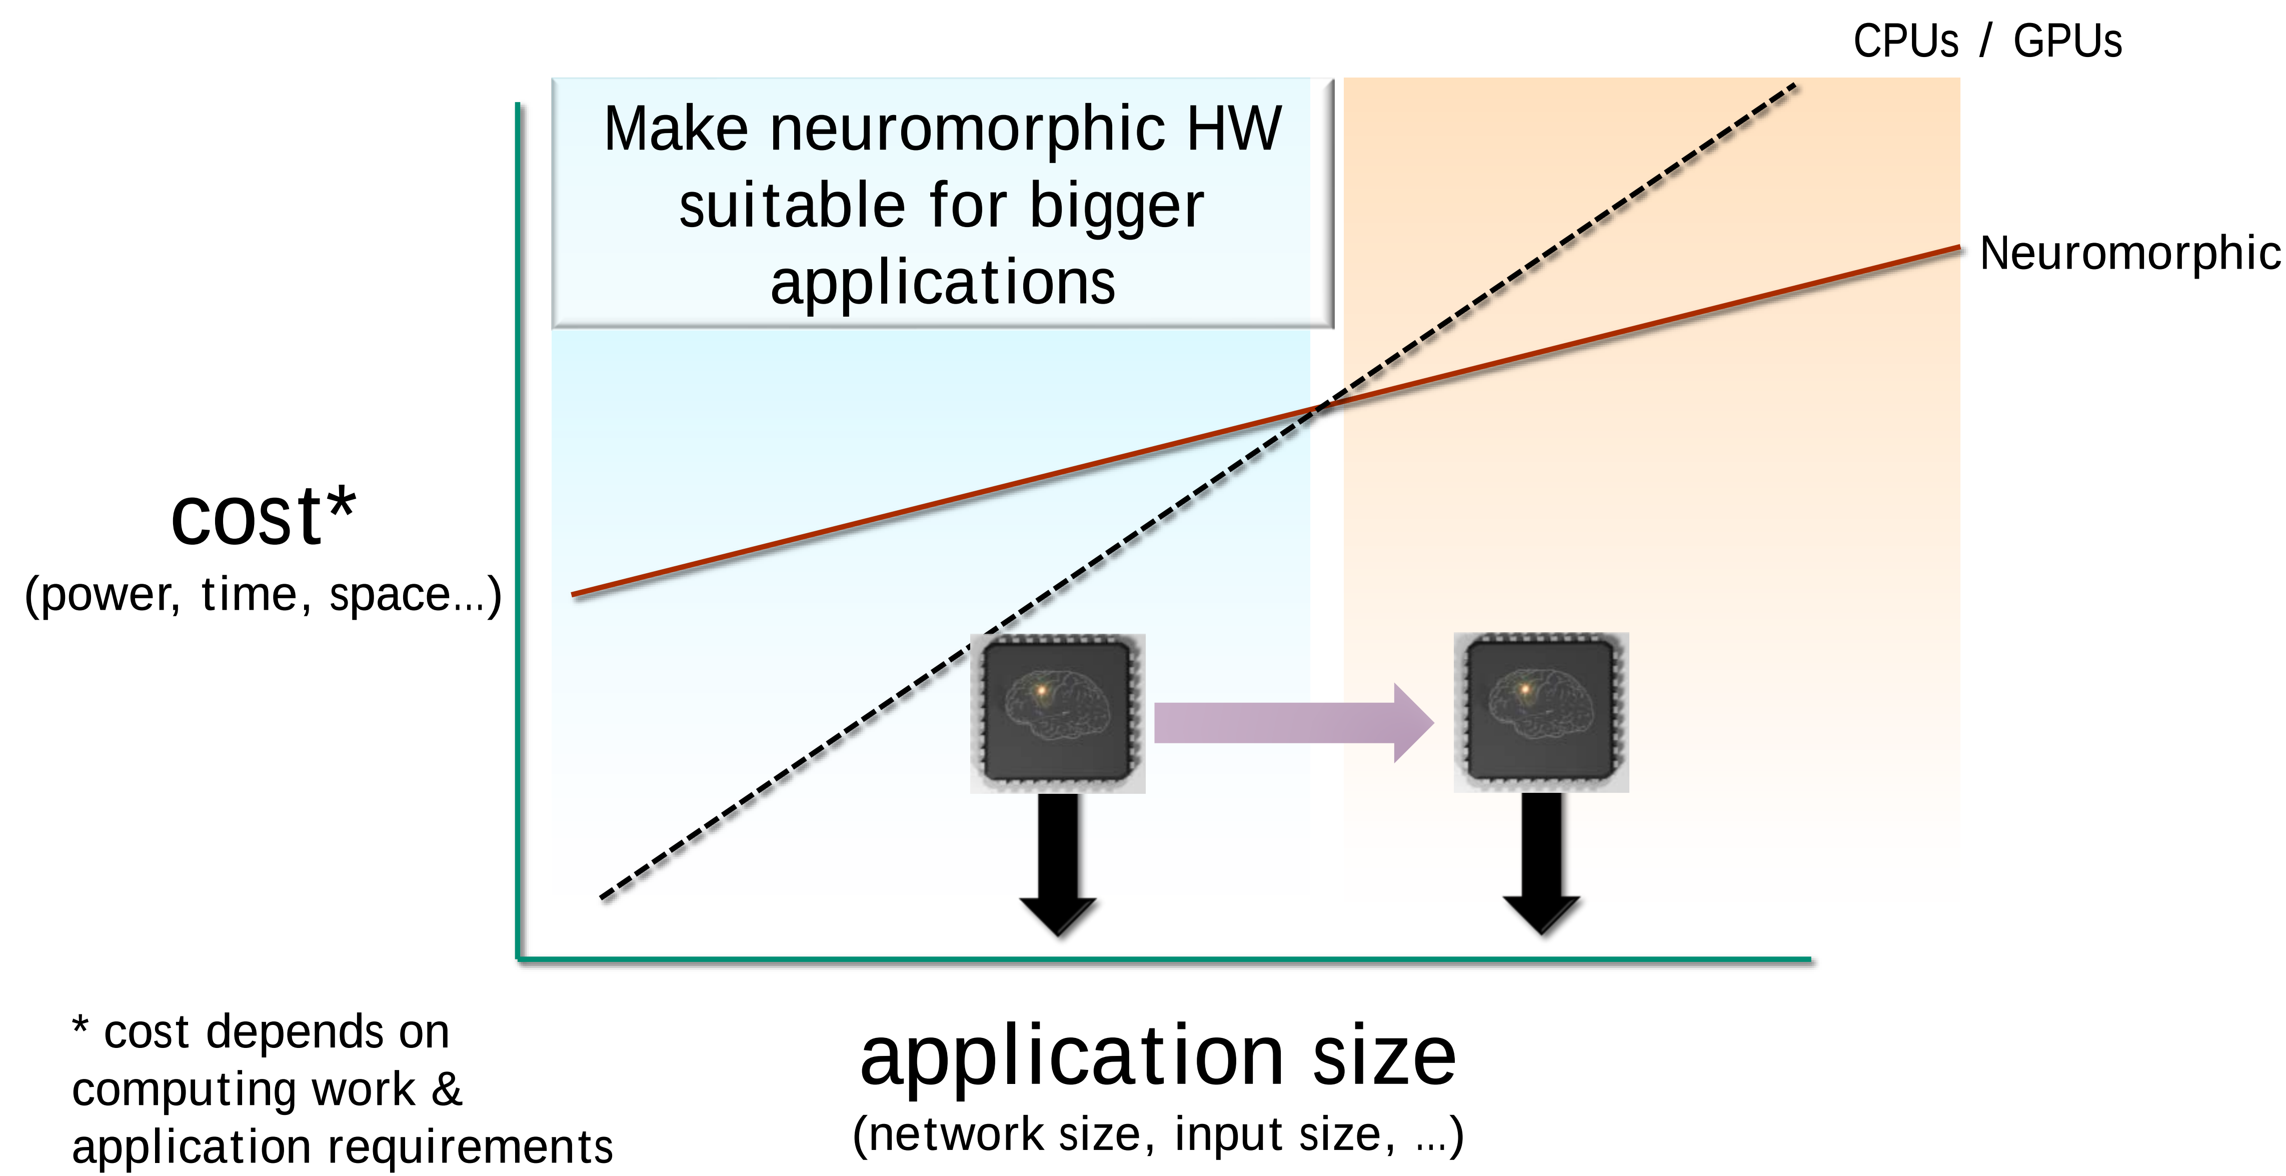

* cost depends on computing work & application requirements

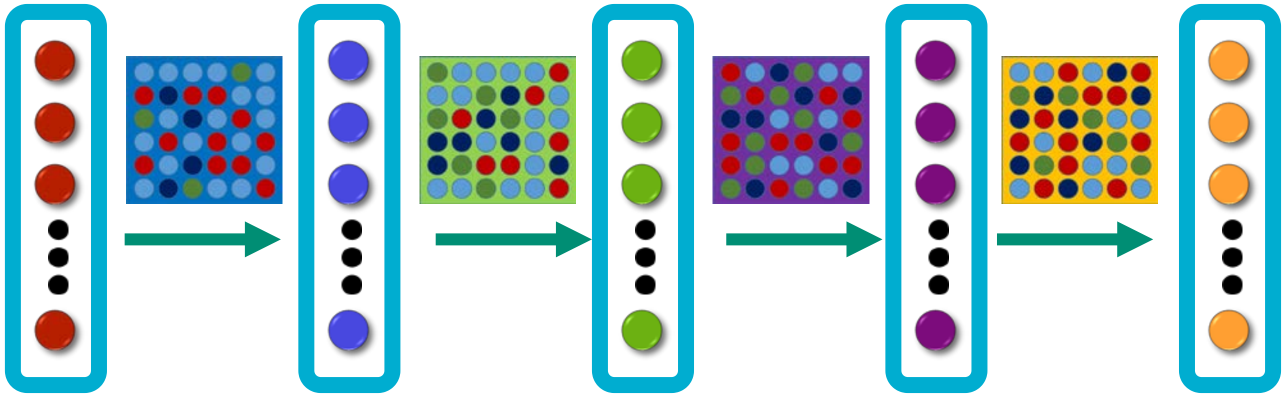

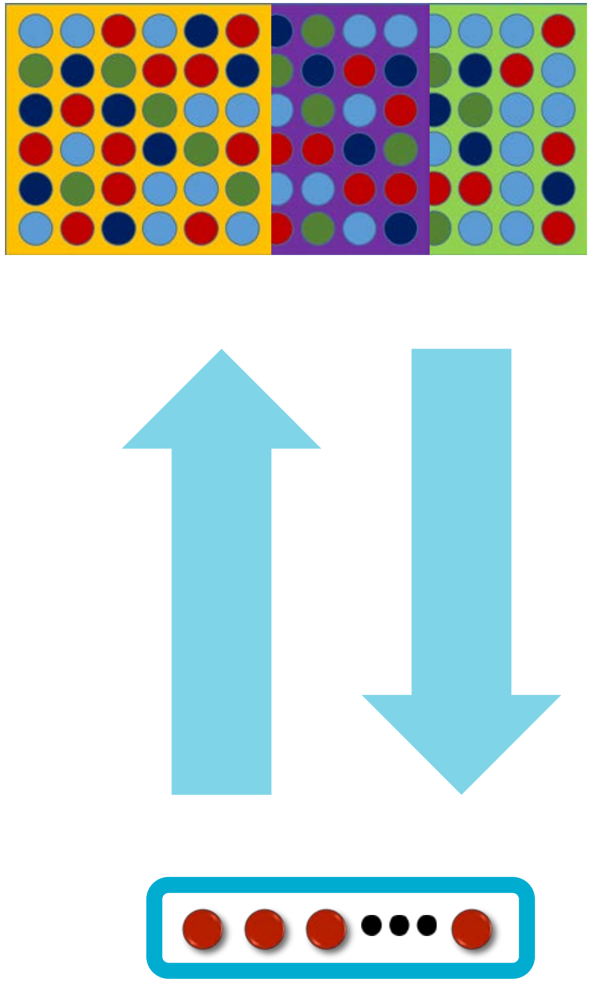

Memory is cheap. We should take advantage of that.

More synapses (whether as arrays or tables or whatever) is not really the problem

Communication is expensive. And it scales poorly with neurons and synapses…

Can we formulate an approach to mitigate this?

Neurons are not as cheap. And we need a lot of them.

But depending on the algorithm, they are often unused for large periods of time.

Can we better take advantage of what we have?

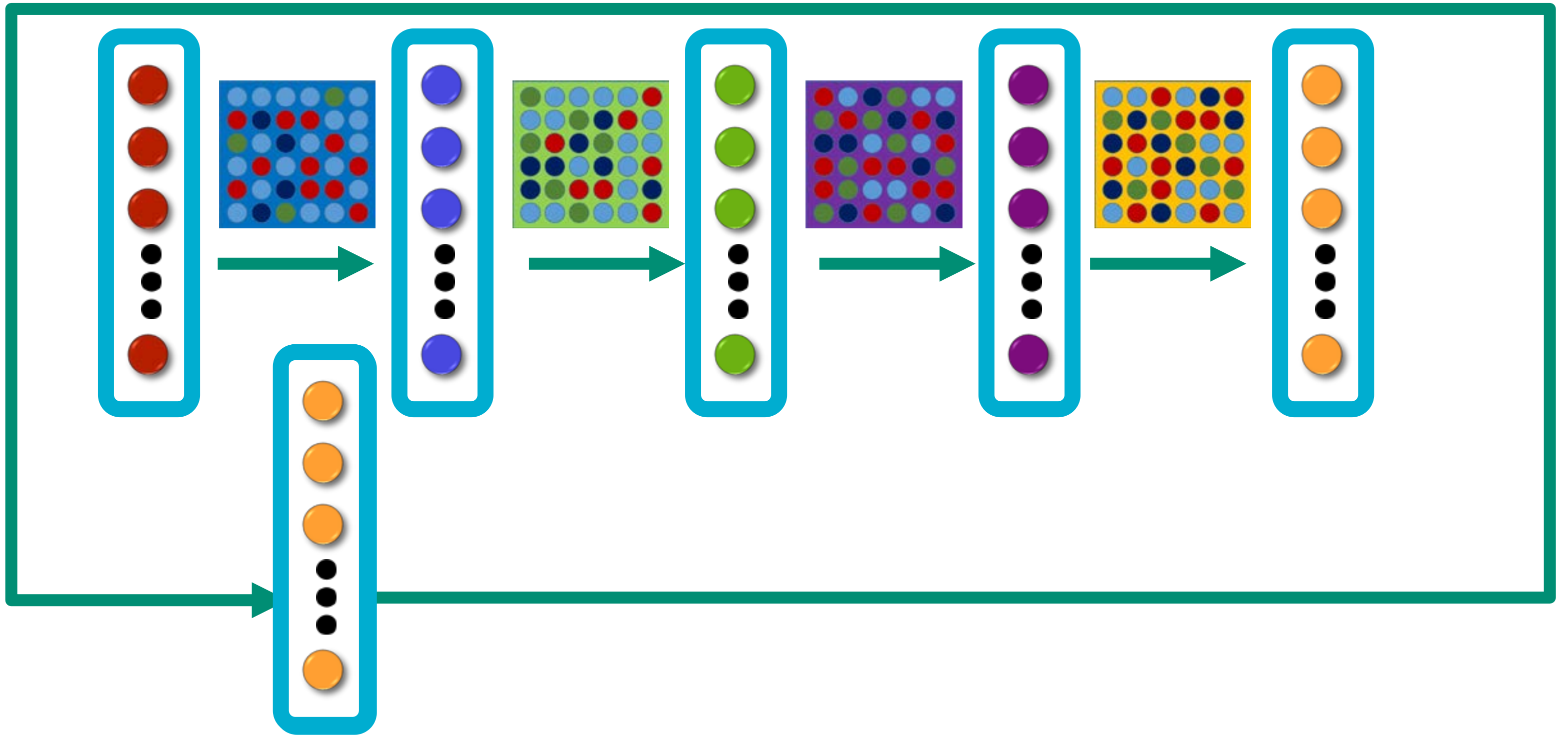

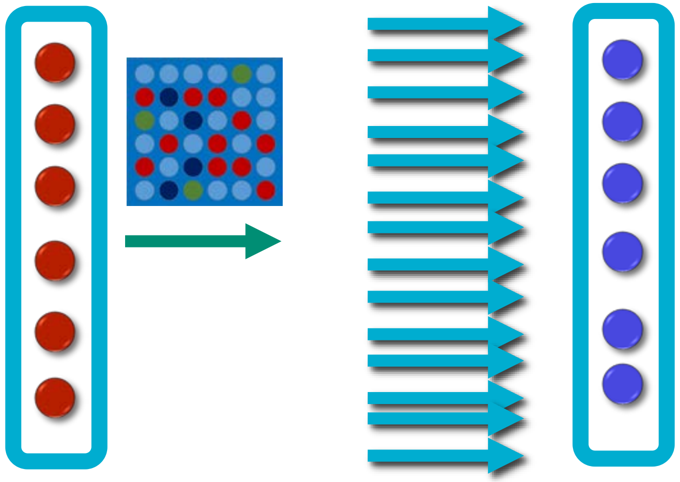

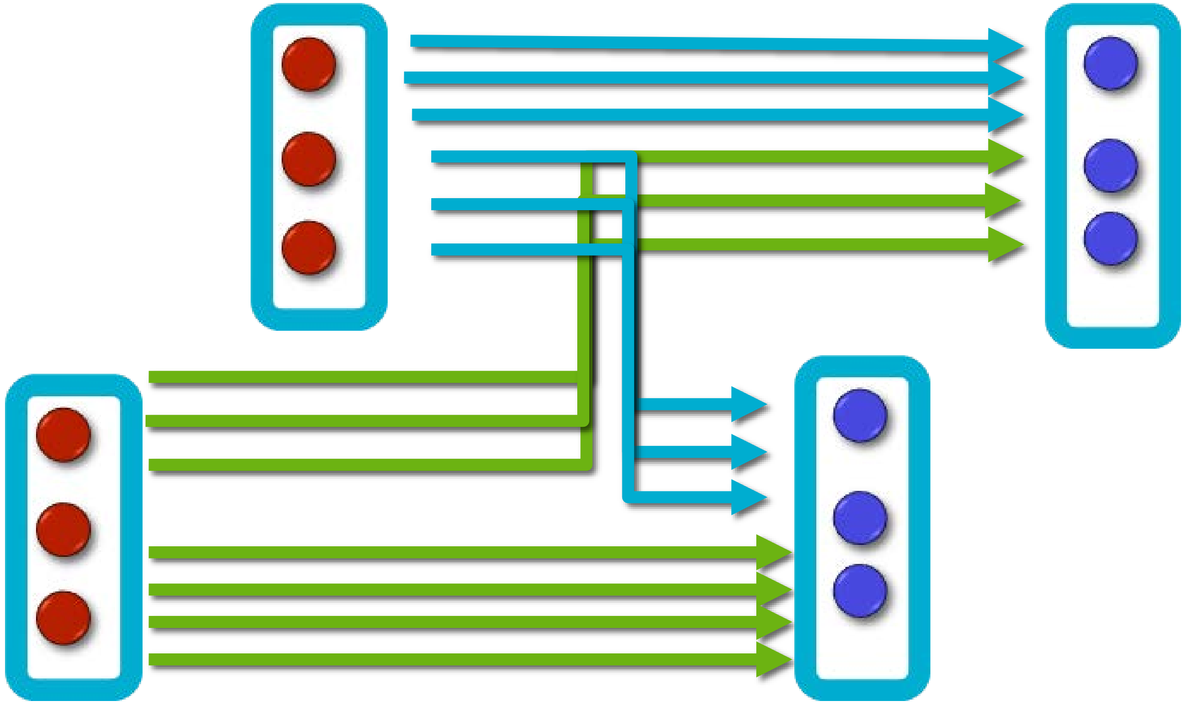

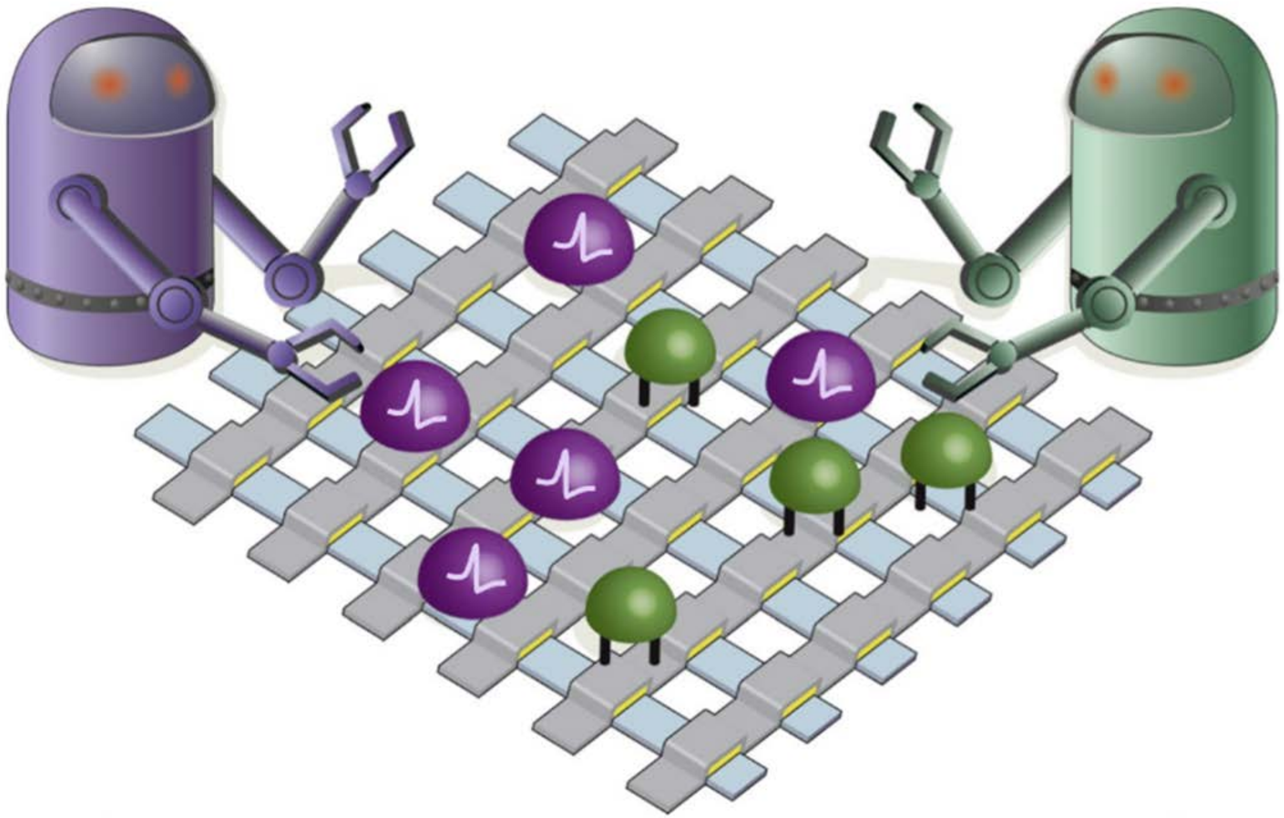

Green & Aimone, Nature Electronics 2019

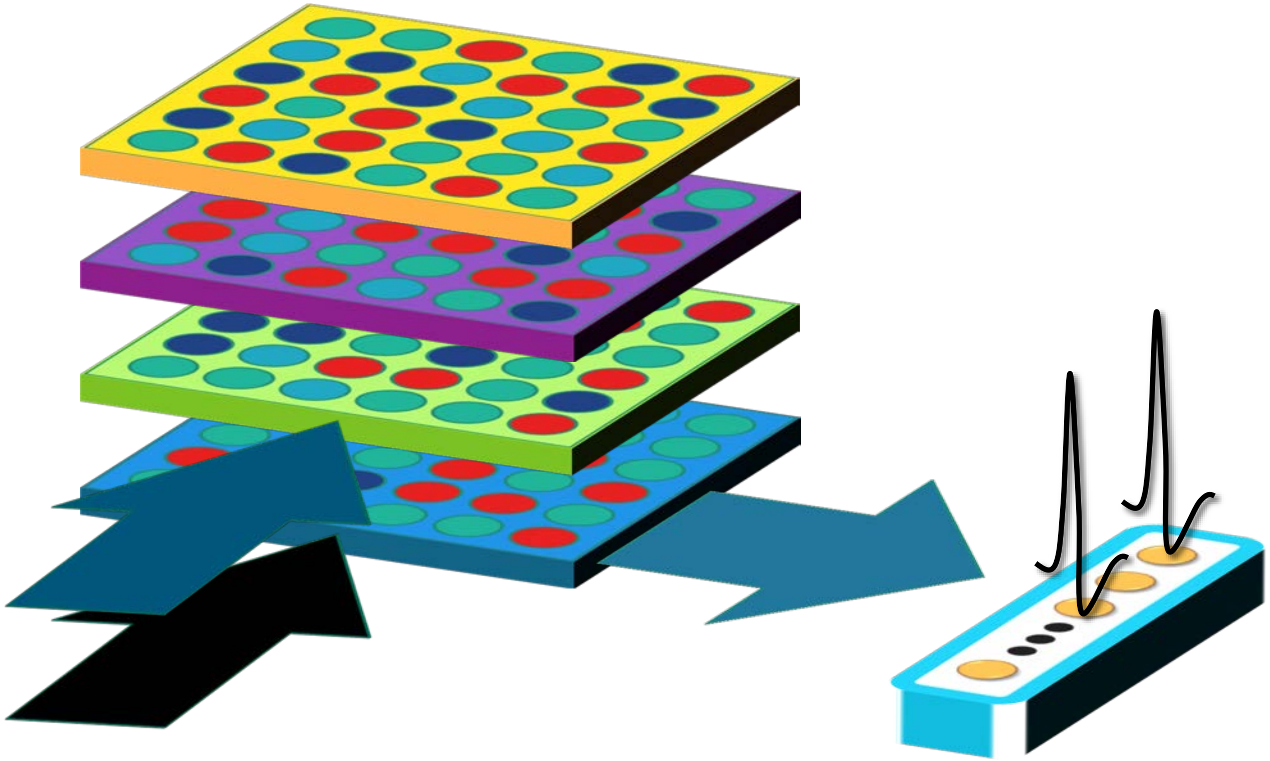

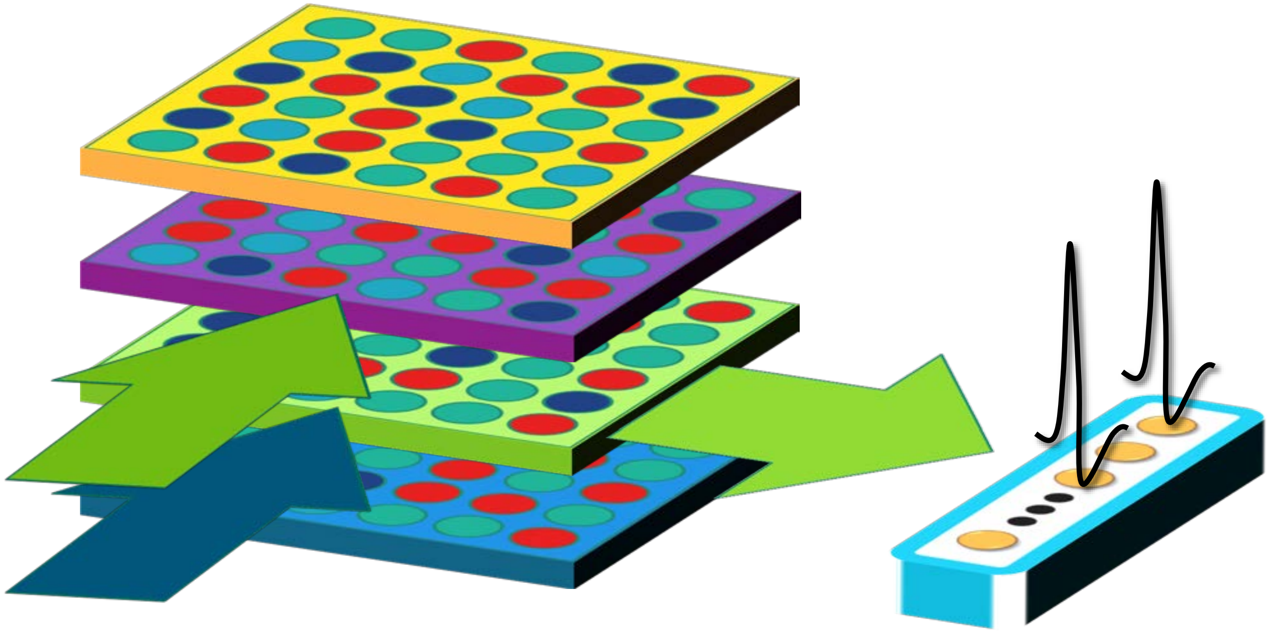

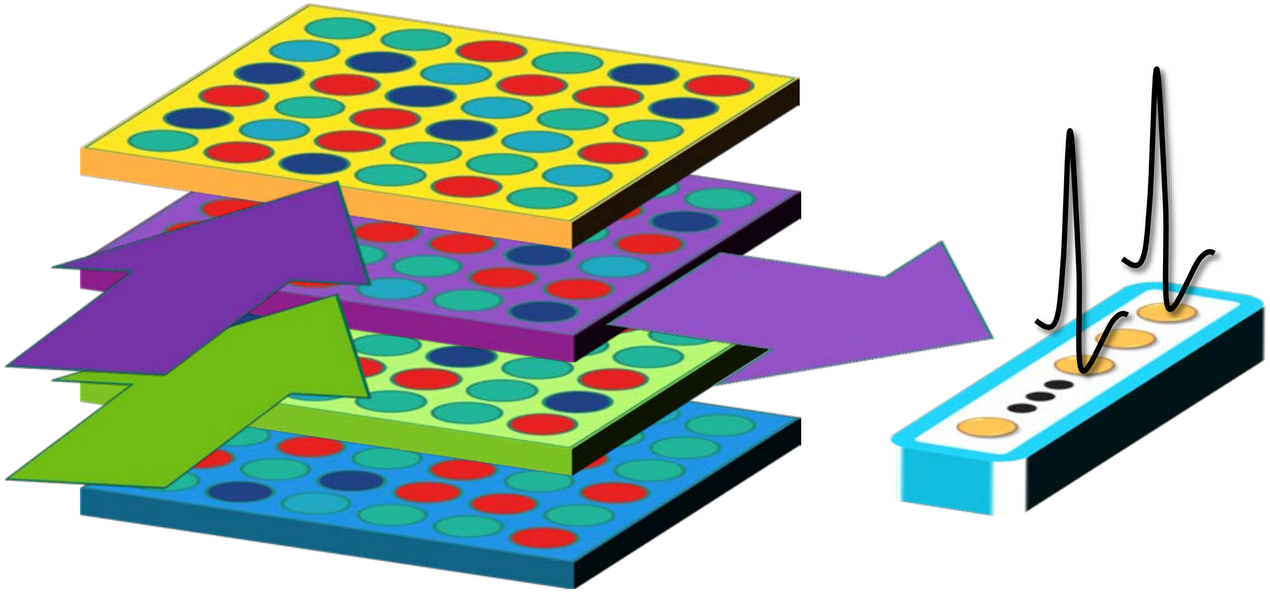

Obviously this won’t work for everything

But it will work for a lot of potential applications

Machine Learning

Whetstone

Convolutions

k-Nearest Neighbor

Support Vector Machines

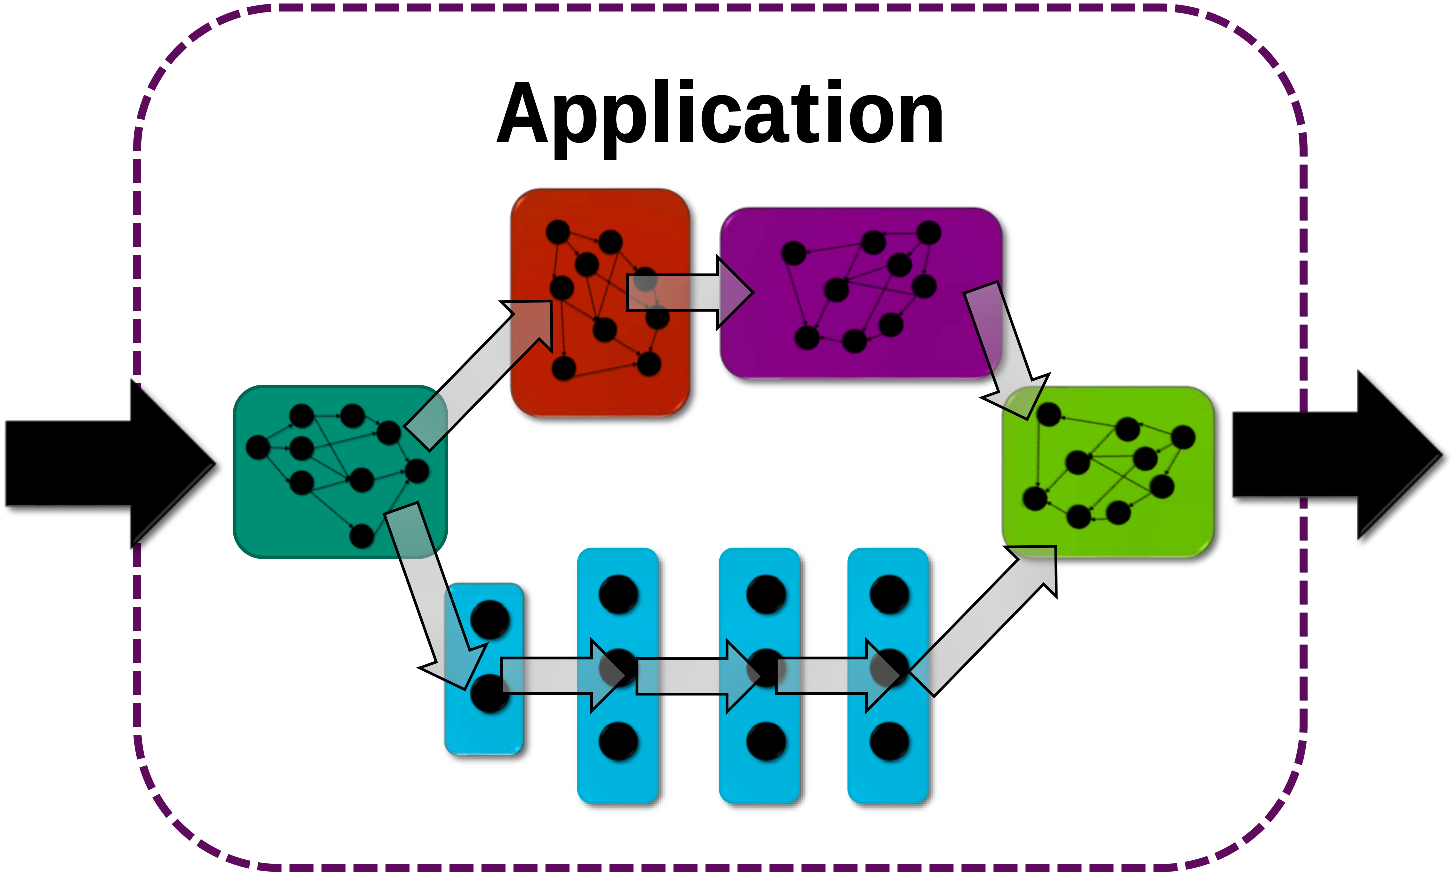

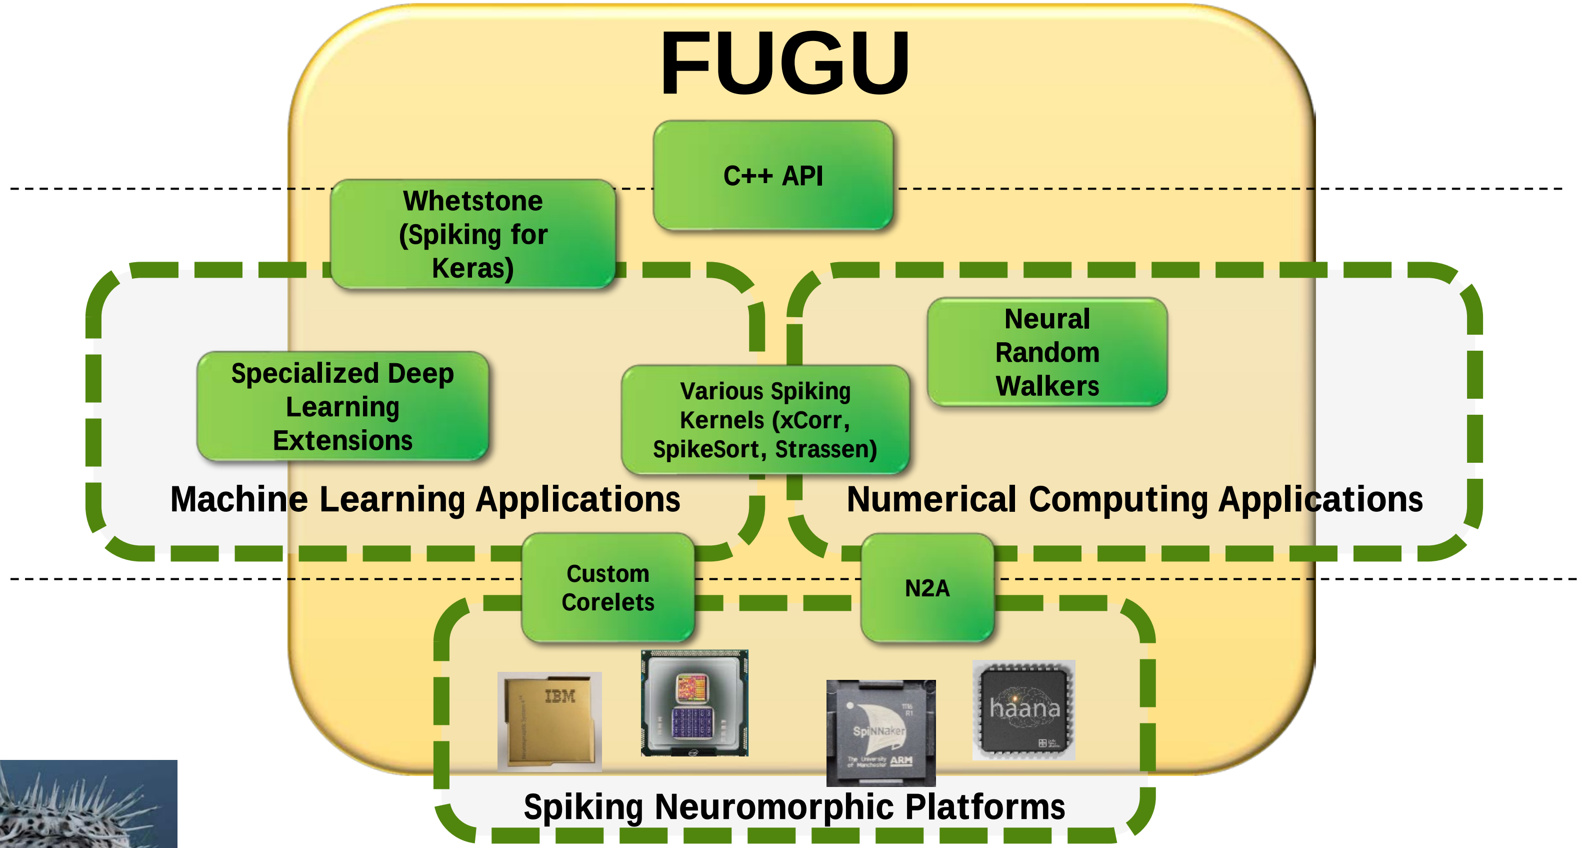

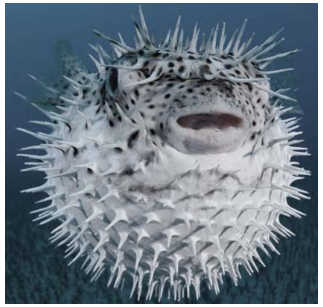

Fugu = pufferfish (why? Pufferfish have spikes…)

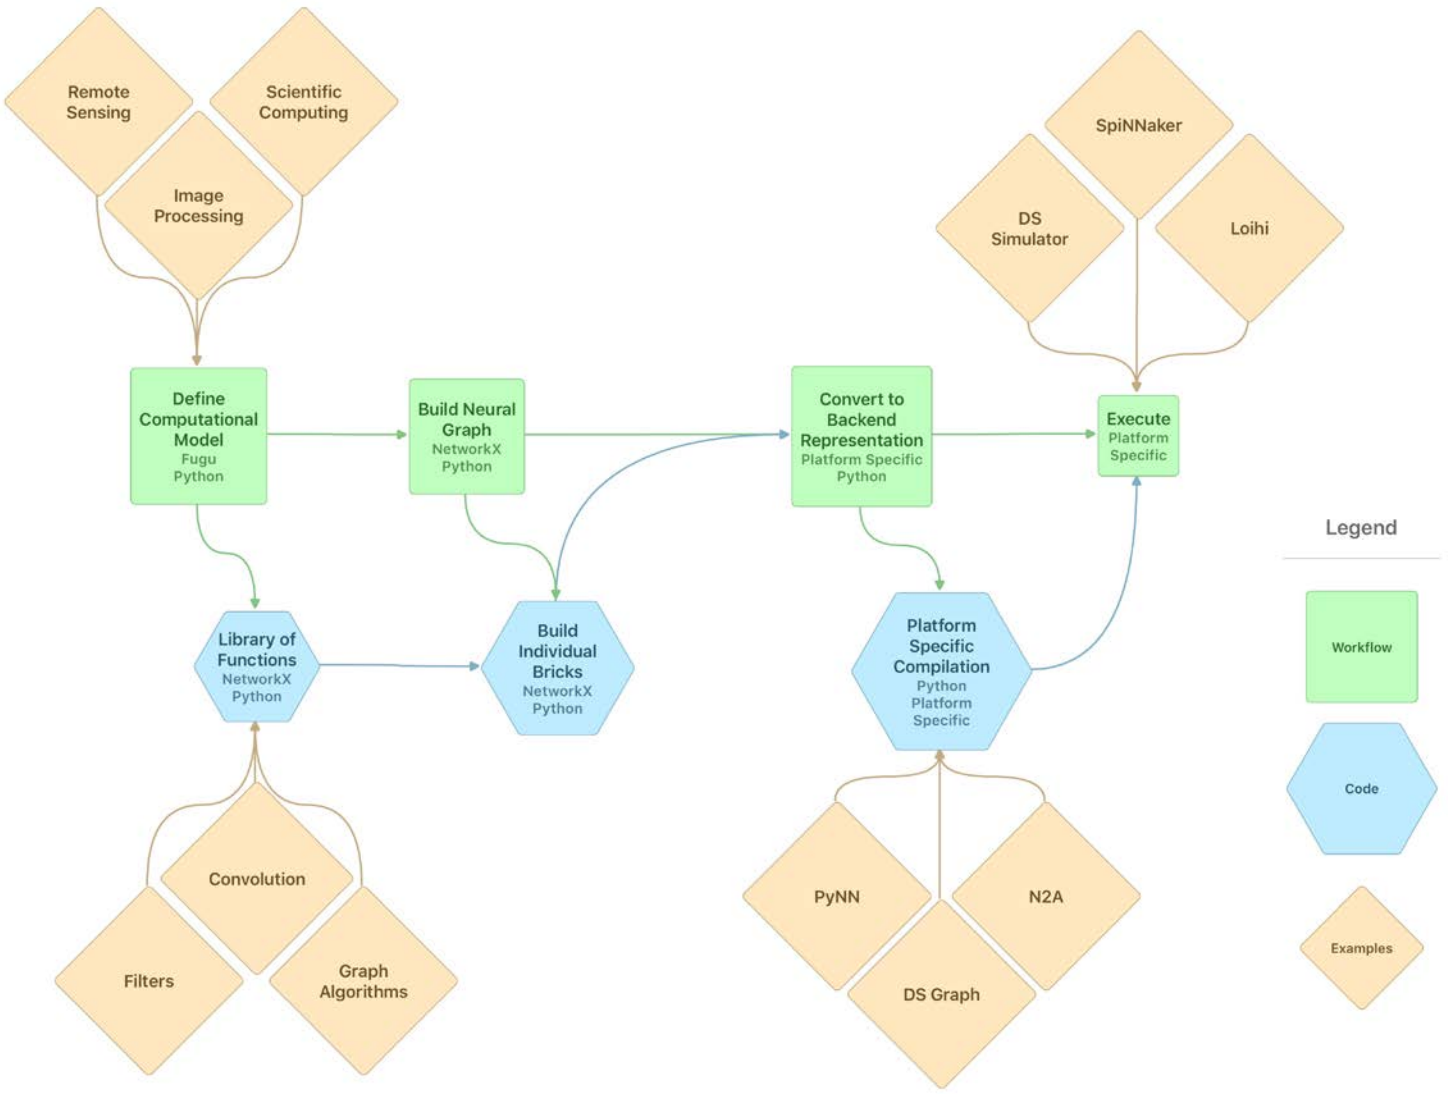

Remote Sensing
Scientific Computing
Image Processing
SpiNNaker
DS Simulator
Loihi
Define Computational Model
Fugu
Python
Build Neural Graph
NetworkX
Python
Convert to Backend Representation
Platform Specific
Python
Execute
Platform Specific
Legend
Library of Functions
NetworkX
Python
Build Individual Bricks
NetworkX
Python
Platform Specific Compilation
Python
Platform Specific
Workflow
Code
Convolution
Filters
Graph Algorithms
PyNN
N2A
DS Graph
Examples

Thanks everyone for coming to NICE 2019!