INTERACTIVE INTELLIGENCE

MOTIVATION

Conscious algorithms, self directed learning, self improvement, accurate computational representations of the world

Why doesn't this exist?

Narrow AI is more profitable

Image classification, stock market prediction, facial recognition, and NLP based personal assistants are all far more lucrative than research

Our algorithms are still wrong

The famed AlphaGo algorithm which beat Lee Sedol at chess cannot answer the question "what is Go?" -- it does not represent the game it has mastered in the context of the world

Computational architecture is still wrong

OpenAI's GPT-3, an NLP model capable of near human conversation took around 190,000 kWh of power to train. If a human brain used that much energy, we would be 1100 years old before we learned how to talk (poorly at that)

MACHINE LEARNING - THE HUMAN WAY

Fundamentally, humans grow up learning through interaction with their environment. Thus, so will our approaches.

A network that can accurately classify images of dogs or cats does not know what a cat is -- it has no idea that a cat purrs or is inherently evil and secretly plotting world domination because the convolutional network has only interacted with images of cats, not cats.

RIP Pepper

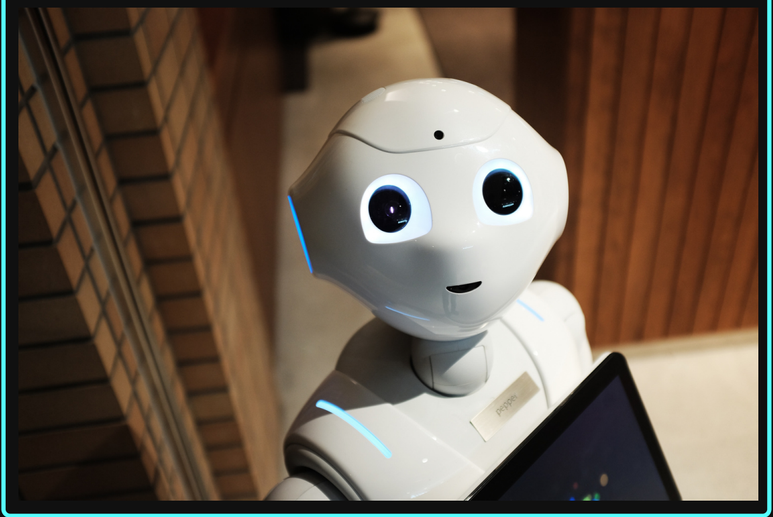

TEAM GOALS

1. Innovate our own, biologically inspired approaches to implement computational learning
2. Implement and test our ideas
3. Examine the differences between computational and biological intelligence
4. Learn, experiment, bring together neuroscience and CS

FIRST PRINCIPLES

1. Human level machine intelligence is not solved
2. Currently, many diverse approaches are being tried from distributed networks of AI to neuromorphic computing
3. Your ideas could contribute
4. We're gonna try them!

LEADERSHIP

Janna Hong
Neuroscience Lead

Chaytan Inman
Team Captain

Chaytan Inman
ML Lead

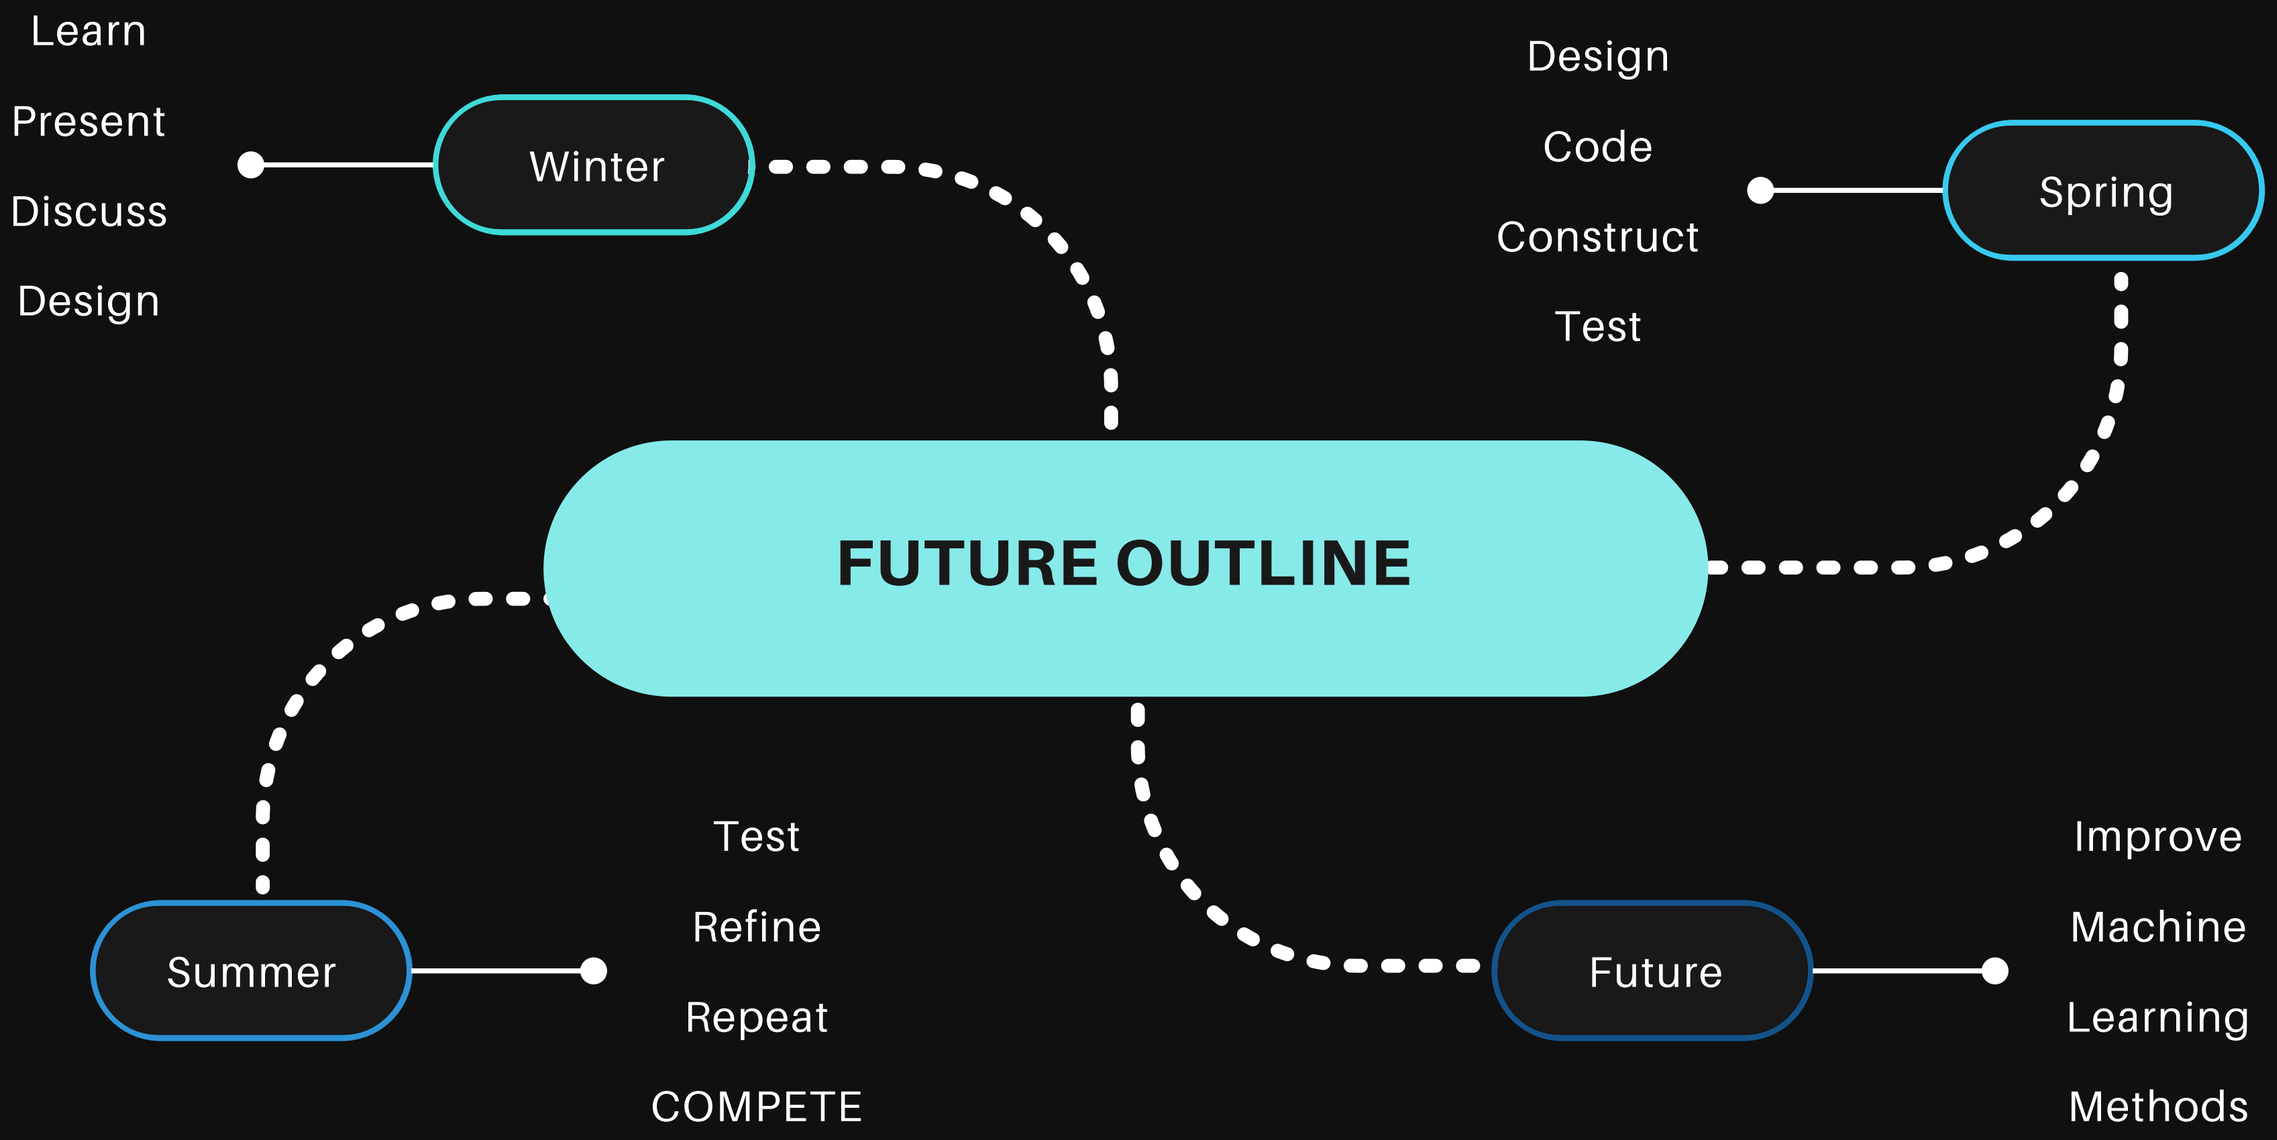
Learn
Present
Discuss
Design
Winter
Design
Code
Construct
Test
Spring
FUTURE OUTLINE
Summer
Test
Refine
Repeat
COMPETE
Future
Improve
Machine
Learning
Methods

MEMBER EXPECTATIONS

1. Finish delegated tasks
2. Communicate if you cannot
3. Reach out for help when stuck

Coming Months

Moving forward we will

- brainstorm projects
- research relevant areas of knowledge and consult experts
- form project teams
- implement, experiment, retry
- compete, show, and tell projects at the end of Spring Quarter!!!

Time until competition: 4 months

Ready to start a project?

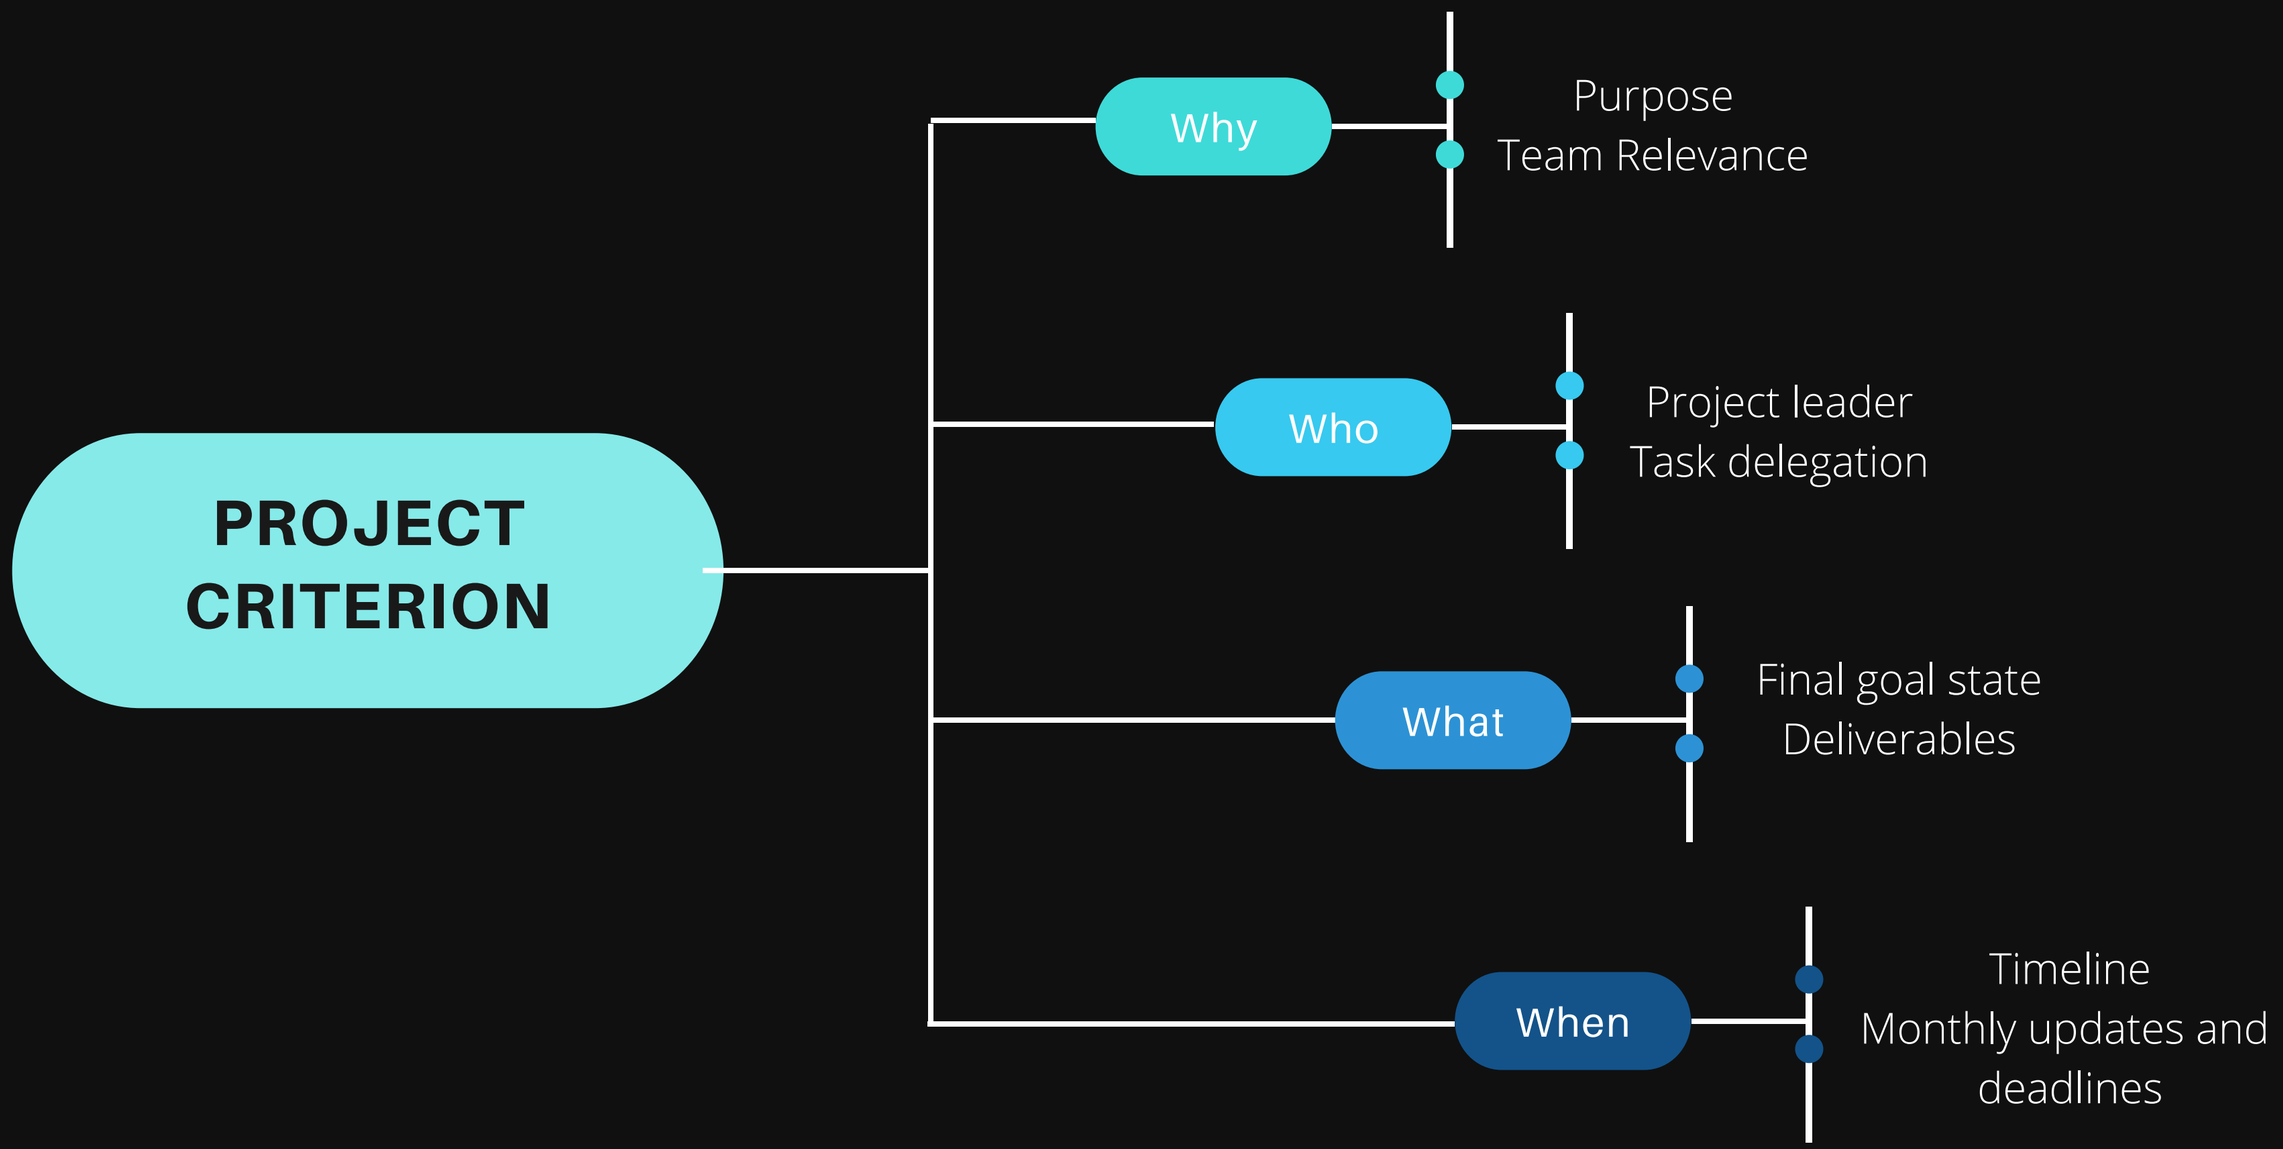
PROJECT CRITERION
Why
Purpose
Team Relevance
Who
Project leader
Task delegation
What
Final goal state
Deliverables
When
Timeline
Monthly updates and deadlines

Project List

Add your idea here

Into the weeds: Artificial General Intelligence

Ongoing Research Areas and Approaches

- Neuromorphic computing
- Deep reinforcement learning
- Neuroscience basis of consciousness
- Brain computer interfaces
- Transfer learning
- Evolutionary and genetic algorithms
- OpenCog platform
- Even deeper deep learning??
- Neural Architecture Search

Ongoing Corporate Attempts

- Microsoft's OpenAI
- Google's DeepMind
- Google Brain
- Facebook AI Research
- Intel Neuromorphic Computing, Loihi

How do you even machine-learn?

Practical Machine Learning

1) Define the problem, environment, and accessible data
2) Choose your racer
Supervised, Unsupervised, Reinforcement
3) Choose your cart
Regression, Classification, Clustering, Policy Optimization
4) Change it up a little bit and write a paper (jk)
5) Play Legos with different architectures

Implementation Frameworks

- Python
- PyTorch
- Tensorflow + Keras

Reinforcement Learning

Key Components

- Action space
- State Space
- Reward Function

Agent observes the state
Agent takes an action that changes(?) the state
Agent receives reward (can be negative)
Agent updates way in which it will act

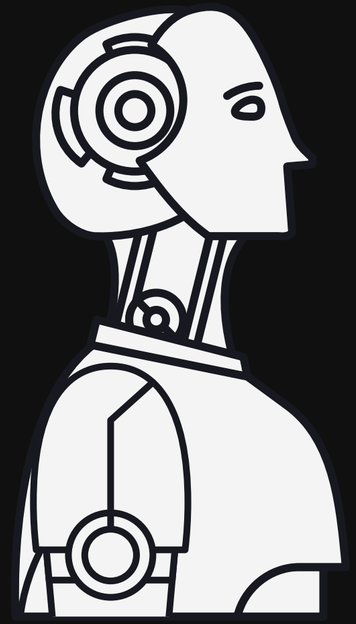

Reinforcement Learning

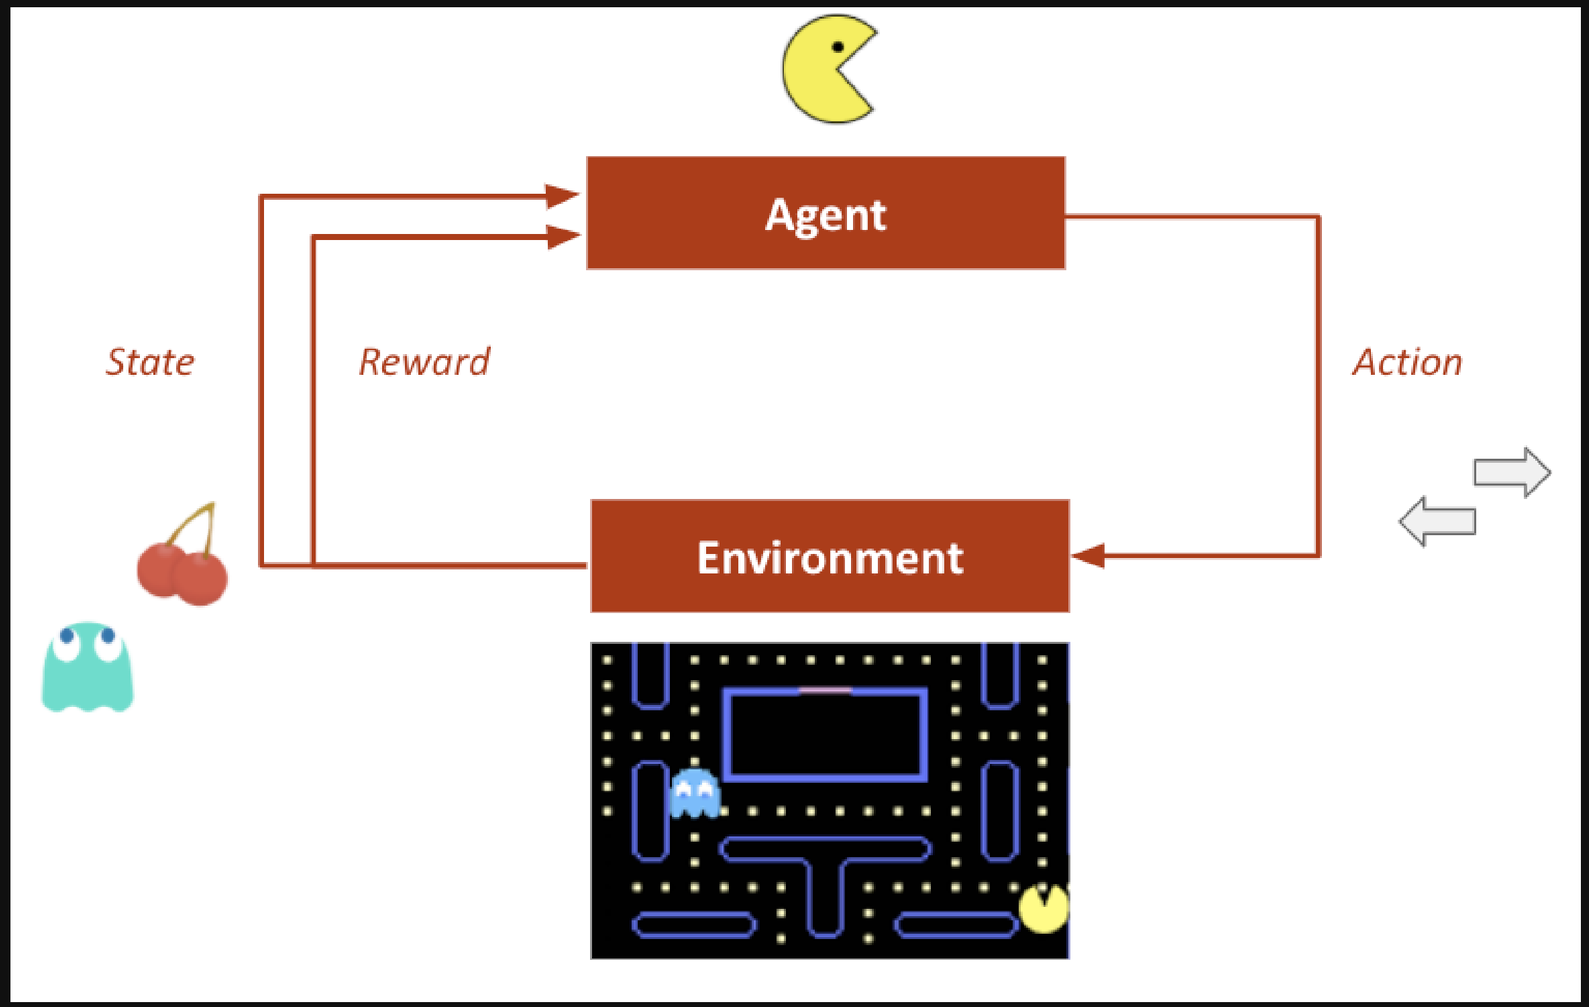

Reinforcement Learning

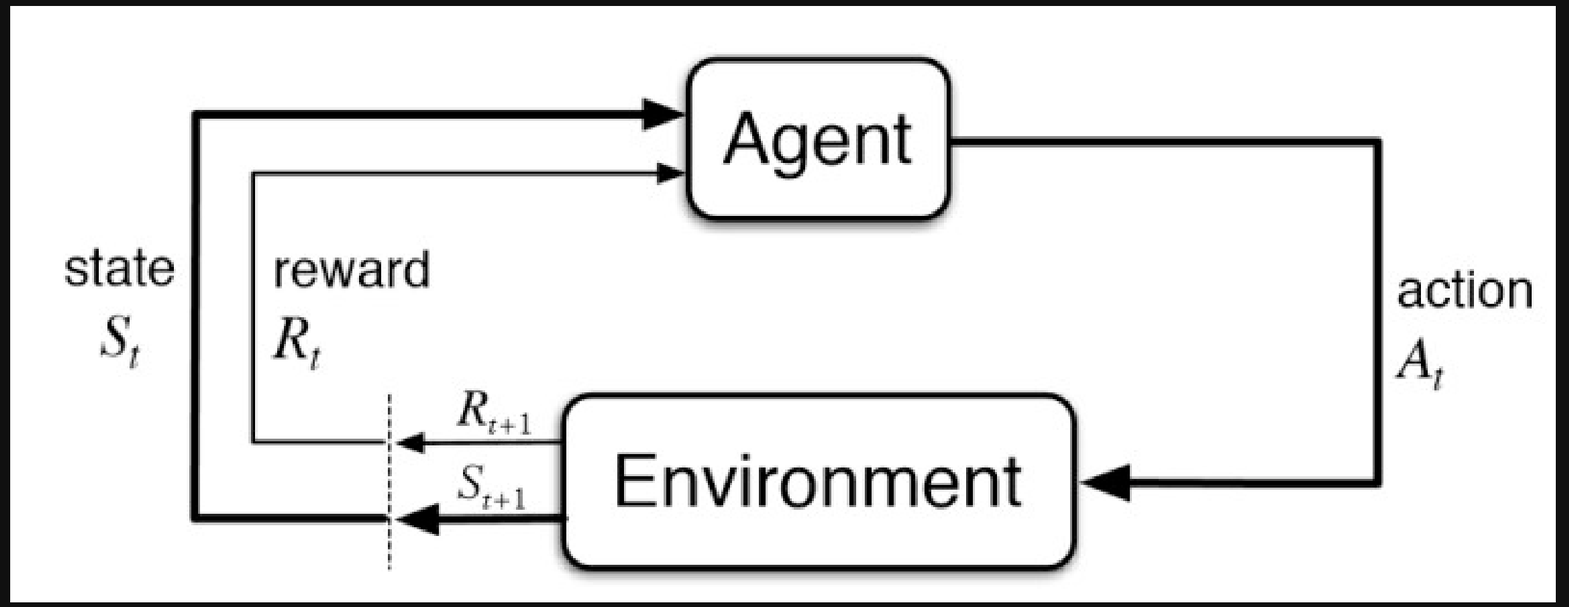

Genetic Algorithms

Key Idea

- Randomize parameters
- Run algorithm
- Check accuracy (fitness function)
- Select some best performers
- Combine their "genetics" / parameters
- Repeat from Step 2 until converging at decent accuracy

** Evolutionary Algorithms are similar but agents die off instead of being evaluated and selected by a fitness function

Neural Architecture Search

Key Ideas

- Use various algorithms (often ML) to find optimal ML architectures
- Often paired with reinforcement learning and RNNs
- (Could also use evolutionary algs or others!)

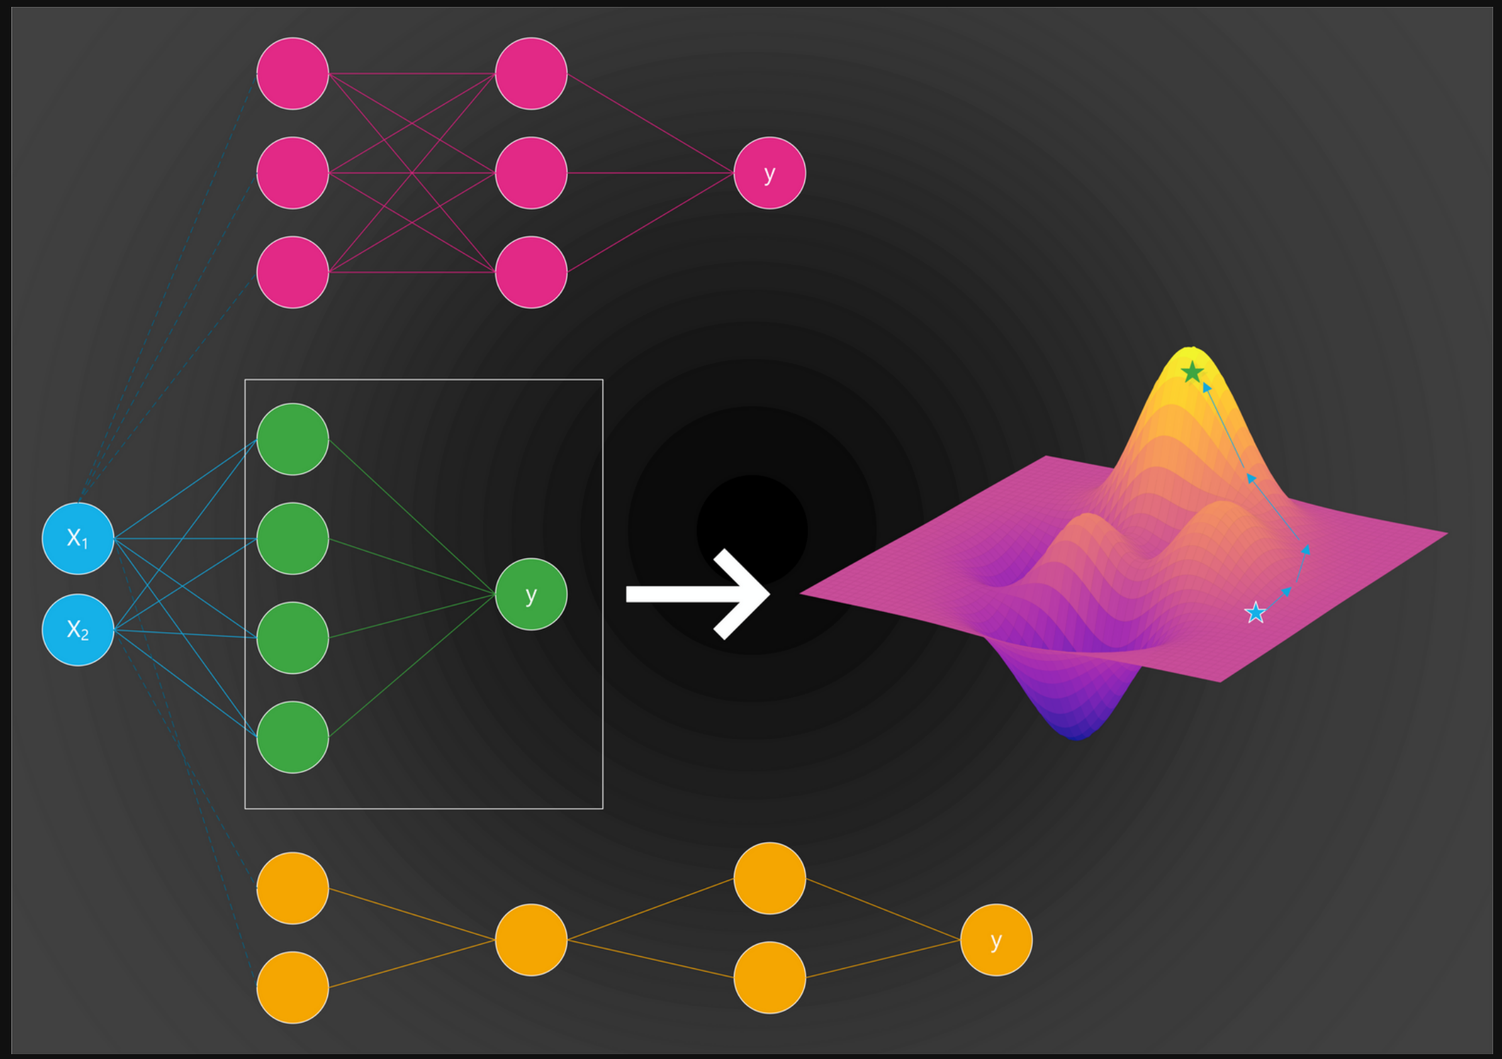

Neuromorphic Computing

Key Idea

- Computation and state should not be separate
- MANY approaches
  - Software (neural networks are states that encapsulate a computation)
  - Hardware (neuromorphic chips)

Different approaches exist within hardware, from mixes of classical Von Neumann architecture to radical biological methods (organoids)

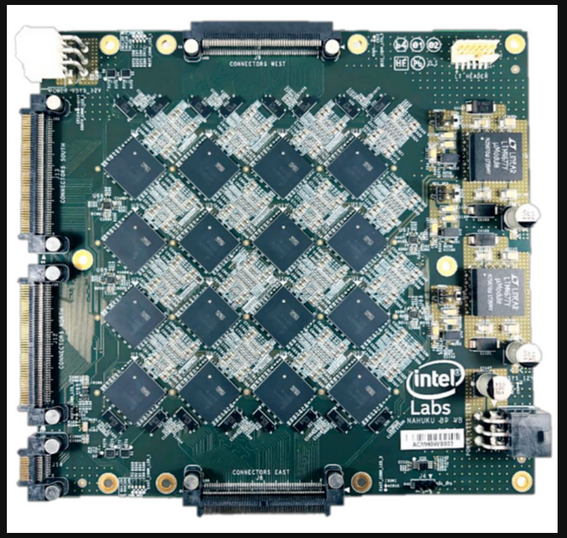

OpenCog Platform

- Hypergraph representation
  - (please do not inquire further about what that means)

- Different modules (reminiscent of different processing areas of the brain ie hippocampus, amygdala etc)

- Re-engineer high level cognitive functions working together

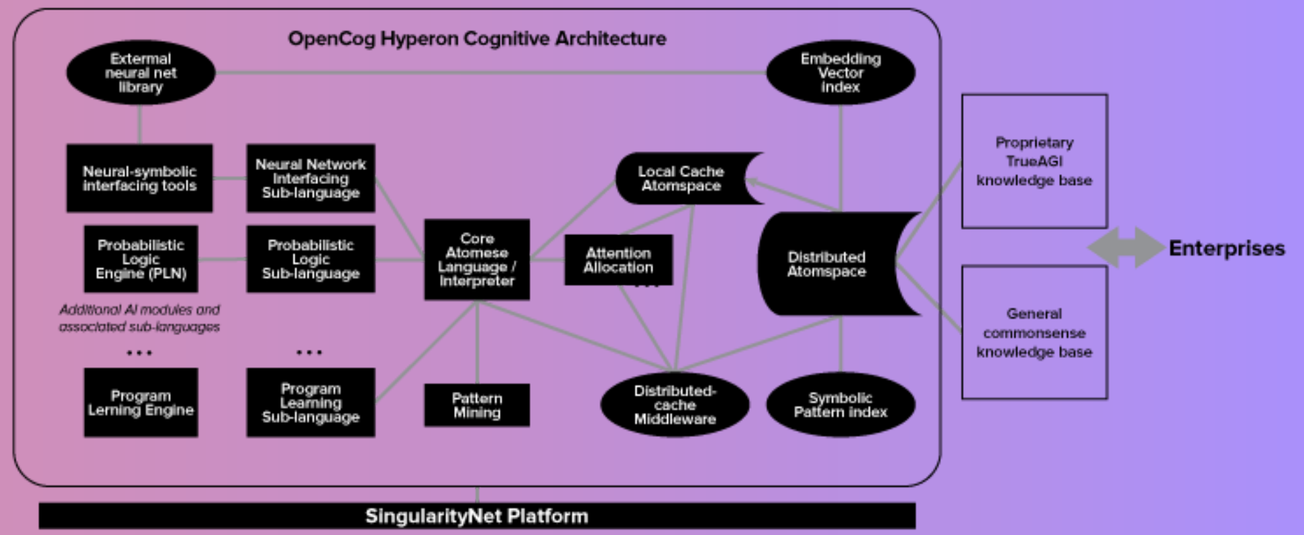

OpenCog Platform

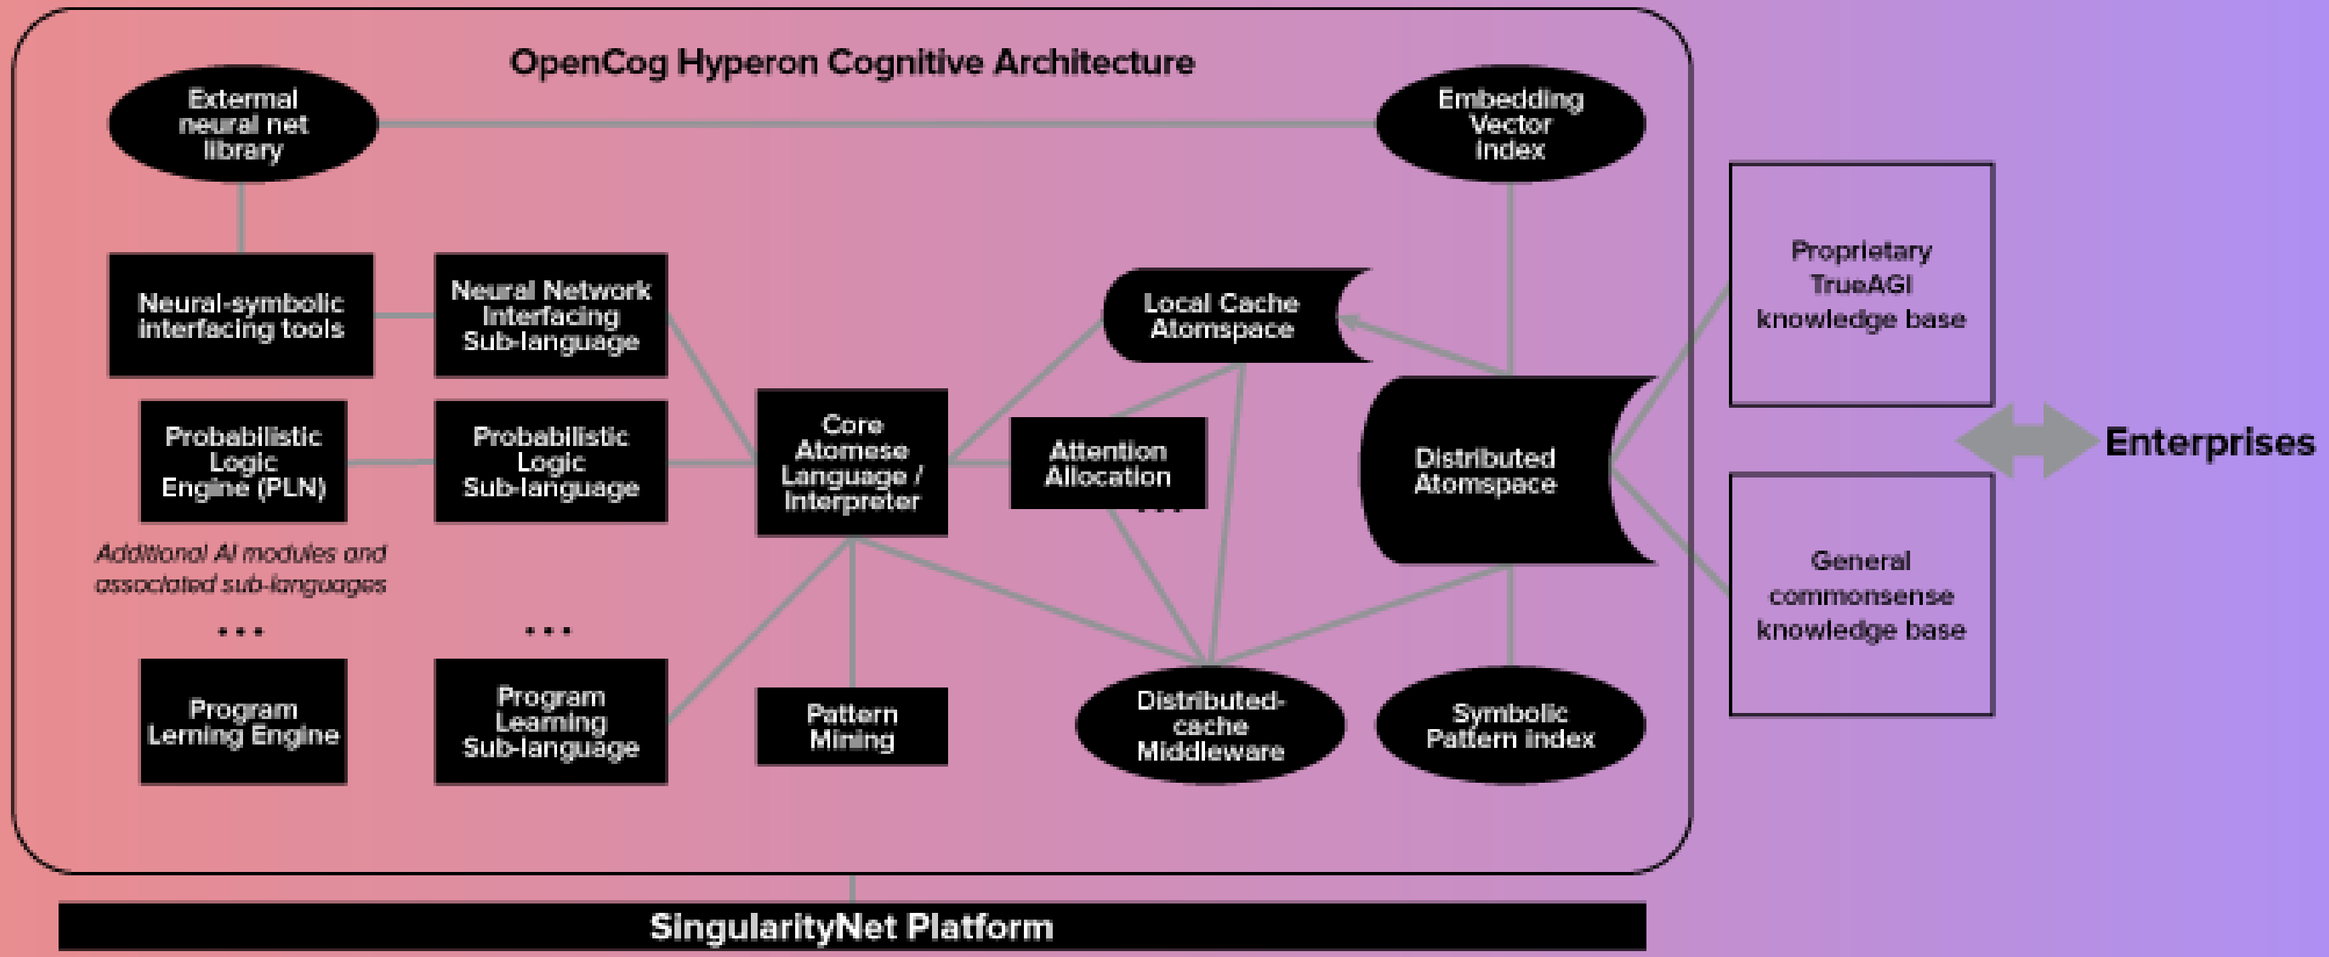

AGI Experts (that I know of)

Ben Goertzel
OpenCog,
SingularityNet

Jeff Hawkins
Numenta,
Neuromorphic Computing

Marcus Hutter
AIXI,
Theoretical AGI

There are a lot of ways to try to engineer more intelligent systems

But what even is intelligence?