

# Ray

## Distributed Hyperparam Tuning

Ray Core  
Ray Tune  
Project  
ASHA, PBT

# Python, Distributed

Run Tasks on all Cores  
(Local, Cluster)

## Ray

- Focus: Training, AI
- Long-Running Actors

## Dask

- Focus: "Pandas" Dataframe
- Arrays, ...

# Ray


- Ray Core
  - Remote Tasks (“function”)
  - Remote Actors (“class”)
  - Logical Resources 8 CPUs, 0.5 GPUs, RAM
- Ray Tune
- Ray Serve, RLlib, ...

# Motivation

- Reinforcement Learning

For Atari, MuJoCo:

- Evolution Strategy also works! (Salimans 2017)




Evolve a Distribution

- ...of NN weights (=Policy)
- Blackbox Optimization
- (sep-)CMA-ES

vs. PPO, GA, NEAT, CEM, ...

# Simulation



<https://log2.ch/progenitor/build15>

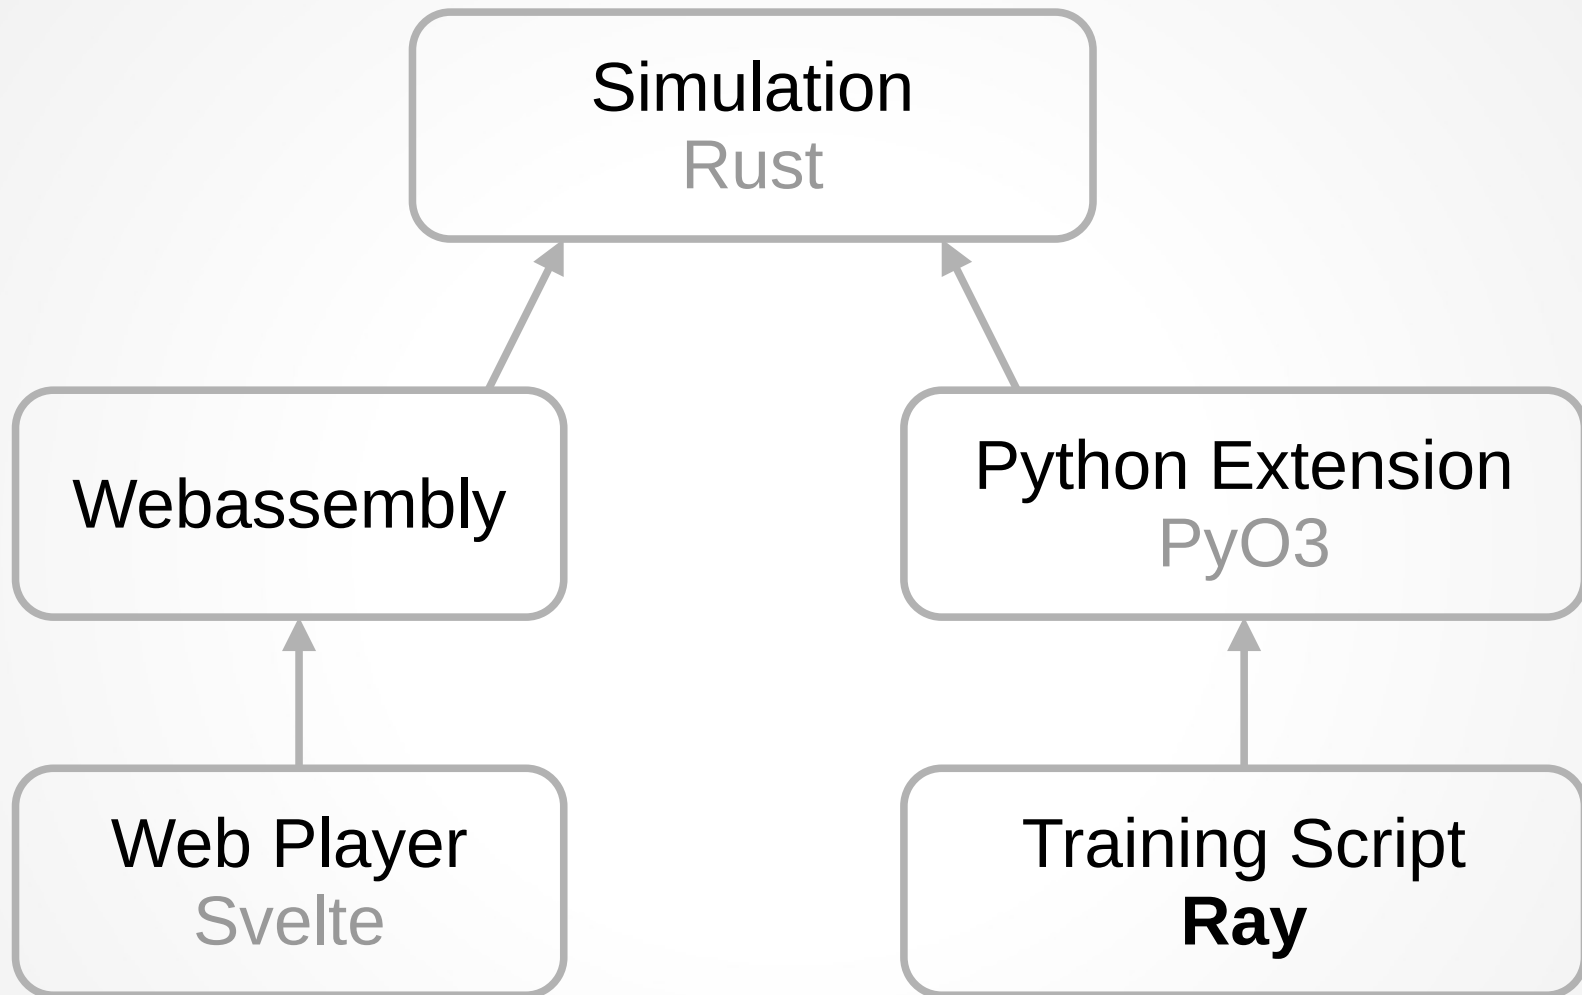
## Collect Food

- Global Reward
- $(\text{Food Neighbours})^2$

## Challenges

- Dig
- Don't Block Others
- Don't Bury Self

# Simulation



# Parameters

```
from ray import tune

param_space = {
    "optimizer": tune.choice(["cmaes", "sep-cmaes"]),
    "popsize": tune.lograndint(30, 200),
    "seeding": tune.choice(["random", "epoch"]),
    "episodes_per_eval": 60,
    "rust": {
        "actions_scale": tune.loguniform(2.0, 40.),
        "bias_fac": 0.1,
        "init_fac": tune.loguniform(0.05, 0.8),
        "n_hidden": 20,
        ...
    }
}
```

# Trainable

```
class MyTrainable(ray.tune.Trainable):
    @classmethod
    def default_resource_request(cls, config):
        return ray.tune.PlacementGroupFactory(
            [{'CPU': 0.01}] + [{'CPU': 1.0}] * 8
        )

    def setup(self, config):
        self.optimizer = cmaes.SepCMA(...)

    def step(self):
        solutions = self.optimizer.ask()
        futures = [evaluate.remote(x, self.config) for x in solutions]
        costs = ray.get(futures)
        self.optimizer.tell(costs)
        return {'score': np.mean(-costs), ... }

    def save_checkpoint(self, tmpdir):
        ...
    def load_checkpoint(self, tmpdir):
        ...
```





### Cluster utilization

**i** Set up Prometheus and Grafana for better Ray Dashboard experience

Time-series charts are hidden because either Prometheus or Grafana server is not detected. Follow [these instructions](#) to set them up and refresh this page.

### Recent jobs

- [gen2-pbt-1 \(06000000\)](#)  
python builders\_optimize.py gen2-pbt-1 --storage=s3://r
- [gen2-asha-1 \(05000000\)](#)  
python builders\_optimize.py gen2-asha-1 --storage=s3:/
- [asha-10 \(04000000\)](#)  
python builders\_optimize.py asha-10 --storage=s3://ma
- [asha-9 \(02000000\)](#)  
python builders\_optimize.py asha-9 --storage=s3://max

[View all jobs](#) →

### Serve Deployments

No Deployments yet...

[View all deployments](#) →

## Cluster status and autoscaler

### Node count

**i** Set up Prometheus and Grafana for better Ray Dashboard experience

Time-series charts are hidden because either Prometheus or Grafana server is not detected. Follow [these instructions](#) to set them up and refresh this page.

### Node Status

**Active:**  
7 ray.worker.default  
1 ray.head.default

**Pending:**  
(no pending nodes)

**Recent failures:**  
(no failures)

### Resource Status

**Usage:**  
215.23999999999998/238.0 CPU  
(77.240000000000004 used of 88.24000000000005 reserved in placement groups)  
0B/317.81GiB memory  
0B/130.76GiB object\_store\_memory

**Demands:**  
{'CPU': 1.0}: 1046+ pending tasks/actors (307+ using placement groups)  
{'CPU': 0.01} \* 1, {'CPU': 1.0} \* 11 (PACK): 6+ pending placement groups

## Node Statistics

TOTALx 13

ALIVEx 9

DEADx 4

## Node List

Host  IP  State  Page Size  Sort By  Reverse:  TABLE CARD

< 1 2 >

	Host / Worker Process name	State	ID	IP / PID	Actions	CPU <sup>?</sup>	Memory <sup>?</sup>	GPU <sup>?</sup>	GRAM	Object Store Memory
>	ip-172-31-5-18	ALIVE	d43e9...	172.31.5.18 (Head)	Log	<div style="width: 15.5%;"><div>15.5%</div></div>	<div style="width: 46.6%;"><div>3.55GB/ 7.60GB(46.6%)</div></div>	N/A	N/A	<div style="width: 0.0%;"><div>0.0000B/ 18.37GB(0.0%)</div></div>
>	ip-172-31-39-157	ALIVE	0f855...	172.31.39.157	Log	<div style="width: 99.3%;"><div>99.3%</div></div>	<div style="width: 7.5%;"><div>4.60GB/ 61.67GB(7.5%)</div></div>	N/A	N/A	<div style="width: 0.0%;"><div>0.0000B/ 18.38GB(0.0%)</div></div>
>	ip-172-31-45-24	ALIVE	36913...	172.31.45.24	Log	<div style="width: 99.1%;"><div>99.1%</div></div>	<div style="width: 7.1%;"><div>4.40GB/ 61.67GB(7.1%)</div></div>	N/A	N/A	<div style="width: 0.0%;"><div>0.0000B/ 18.38GB(0.0%)</div></div>
▼	ip-172-31-38-166	ALIVE	644a1...	172.31.38.166	Log	<div style="width: 98.4%;"><div>98.4%</div></div>	<div style="width: 7.5%;"><div>4.62GB/ 61.67GB(7.5%)</div></div>	N/A	N/A	<div style="width: 0.0%;"><div>0.0000B/ 18.38GB(0.0%)</div></div>
	/home/admin/venv/bin/python	ALIVE		2484	Log CPU Flame Graph Stack Trace Memory Profiling	<div style="width: 0%;"><div>0%</div></div>	<div style="width: 0.1%;"><div>73.24MB/ 61.67GB(0.1%)</div></div>	N/A	N/A	N/A
	ray::evaluate_real	ALIVE		2596	Log CPU Flame Graph Stack Trace Memory Profiling	<div style="width: 90.1%;"><div>90.1%</div></div>	<div style="width: 0.2%;"><div>132.41MB/ 61.67GB(0.2%)</div></div>	N/A	N/A	N/A
	ray::evaluate_real	ALIVE		3941	Log CPU Flame Graph Stack Trace Memory Profiling	<div style="width: 86.3%;"><div>86.3%</div></div>	<div style="width: 0.2%;"><div>129.95MB/ 61.67GB(0.2%)</div></div>	N/A	N/A	N/A
	ray::evaluate_real	ALIVE		12741	Log CPU Flame Graph Stack Trace	<div style="width: 85.8%;"><div>85.8%</div></div>	<div style="width: 0.2%;"><div>128.10MB/ 61.67GB(0.2%)</div></div>	N/A	N/A	N/A

# Local vs Cluster

```
.o S3 Snapshot Configuration :-(  
Current best trial: 78858_00007 with score=1.1365739107131958 and params={'episode
```


Trial name	status	iter	total time (s)	score
BuildersTrainable_78858_00011	RUNNING	9	307.194	0.8165
BuildersTrainable_78858_00012	RUNNING	2	69.4126	0.469094
BuildersTrainable_78858_00013	RUNNING	2	73.9114	0.355227
BuildersTrainable_78858_00014	RUNNING	2	67.8321	0.433017
BuildersTrainable_78858_00015	RUNNING	2	71.9439	0.452359
BuildersTrainable_78858_00000	ERROR	10	339.659	1.0381
BuildersTrainable_78858_00001	ERROR	10	400.753	0.480447
BuildersTrainable_78858_00002	ERROR	10	340.328	0.484905
BuildersTrainable_78858_00003	ERROR	10	333.621	0.457657
BuildersTrainable_78858_00004	ERROR	10	329.257	0.676394
BuildersTrainable_78858_00005	ERROR	10	324.968	0.60222
BuildersTrainable_78858_00006	ERROR	10	307.47	0.879825
BuildersTrainable_78858_00007	ERROR	10	328.747	1.13657
BuildersTrainable_78858_00008	ERROR	10	325.828	0.676646
BuildersTrainable_78858_00009	ERROR	10	305.475	0.625767
BuildersTrainable_78858_00010	ERROR	10	323.592	0.726684

```
(BuildersTrainable pid=8583, ip=172.31.47.107)  
2024-04-04 22:16:29,355 ERROR tune_controller.py:1332 -- Trial task failed for tri  
Traceback (most recent call last):  
File "/home/admin/venv/lib/python3.10/site-packages/ray/air/execution/_internal/  
result = ray.get(future)
```

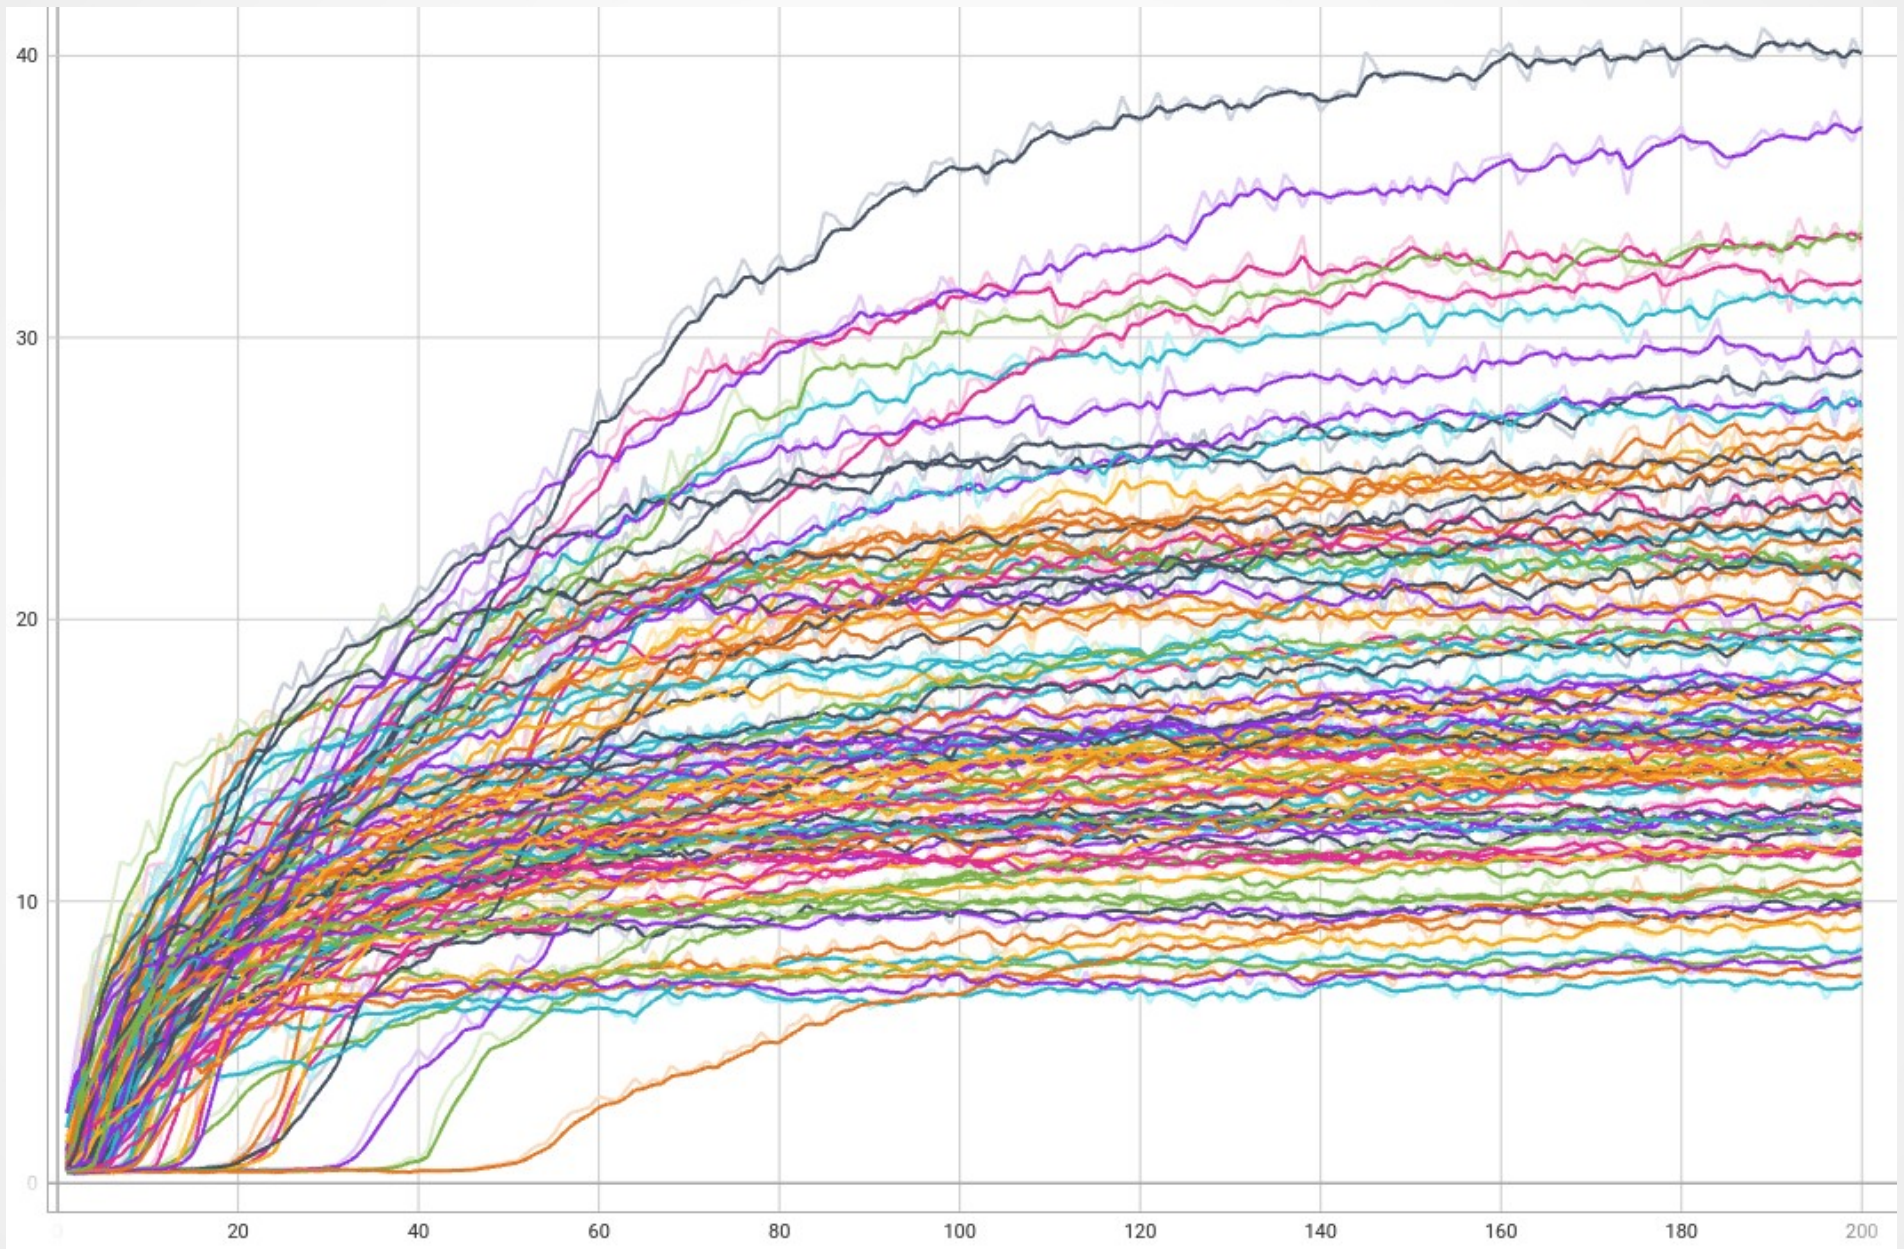
# Random Search

Color by  
ray/tune/score

<b>popsize</b> <input type="radio"/> Linear <input checked="" type="radio"/> Logarithmic <input type="radio"/> Quantile	<b>memory_clamp</b> <input type="radio"/> Linear <input checked="" type="radio"/> Logarithmic <input type="radio"/> Quantile	<b>memory_halftime</b> <input type="radio"/> Linear <input checked="" type="radio"/> Logarithmic <input type="radio"/> Quantile	<b>actions_scale</b> <input type="radio"/> Linear <input checked="" type="radio"/> Logarithmic <input type="radio"/> Quantile	<b>ray/tune/score</b> <input type="radio"/> Linear <input checked="" type="radio"/> Logarithmic <input type="radio"/> Quantile
--	---	--	--	---

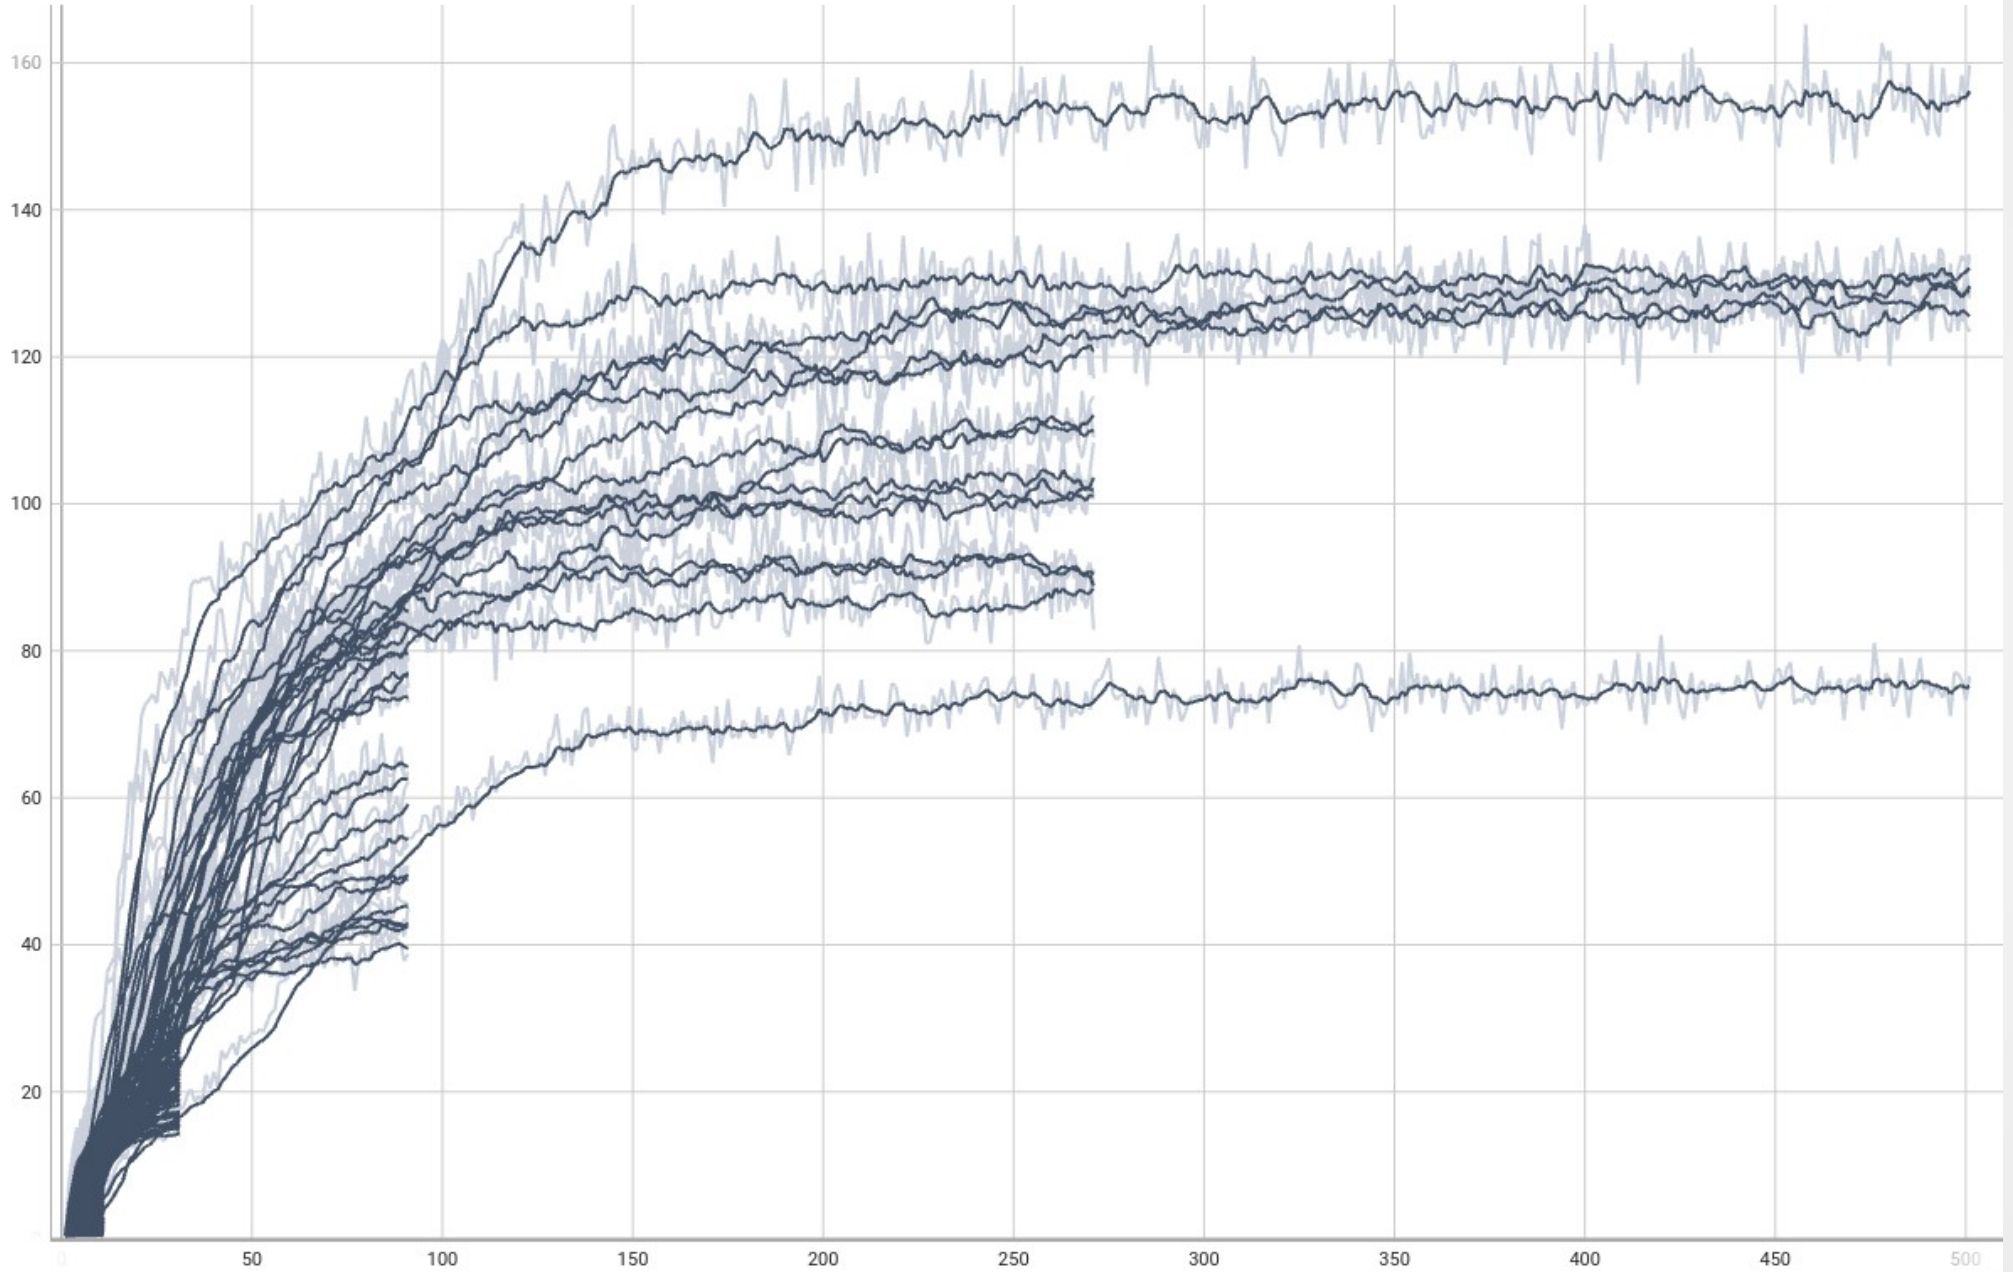


# Random Search

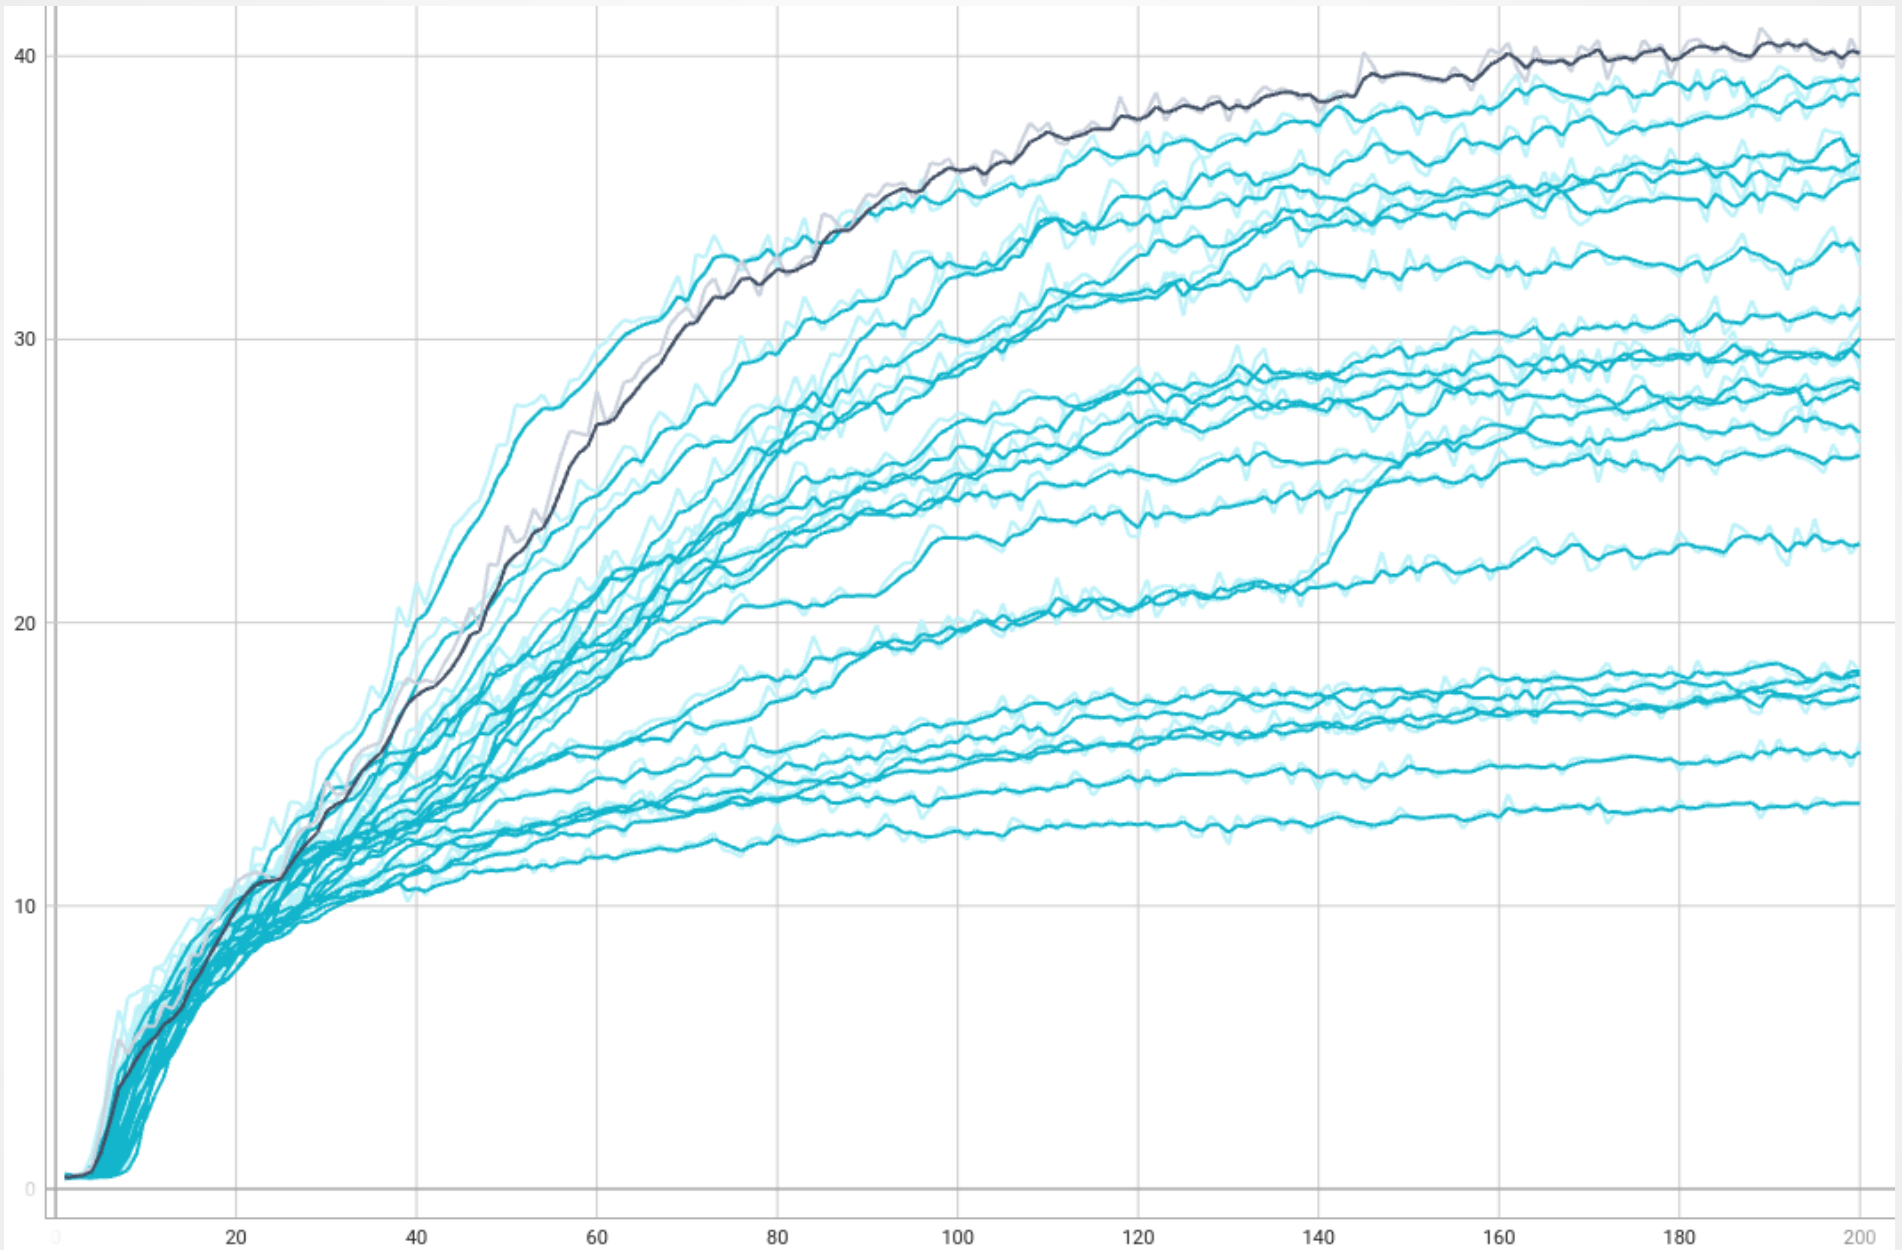


# ASHA

/tune/score



# Reproduce?



# Hyperparams

## ASHA

- “Async Successive Halving”
- **Stop** Bad Trials
- Default Choice

## PBT

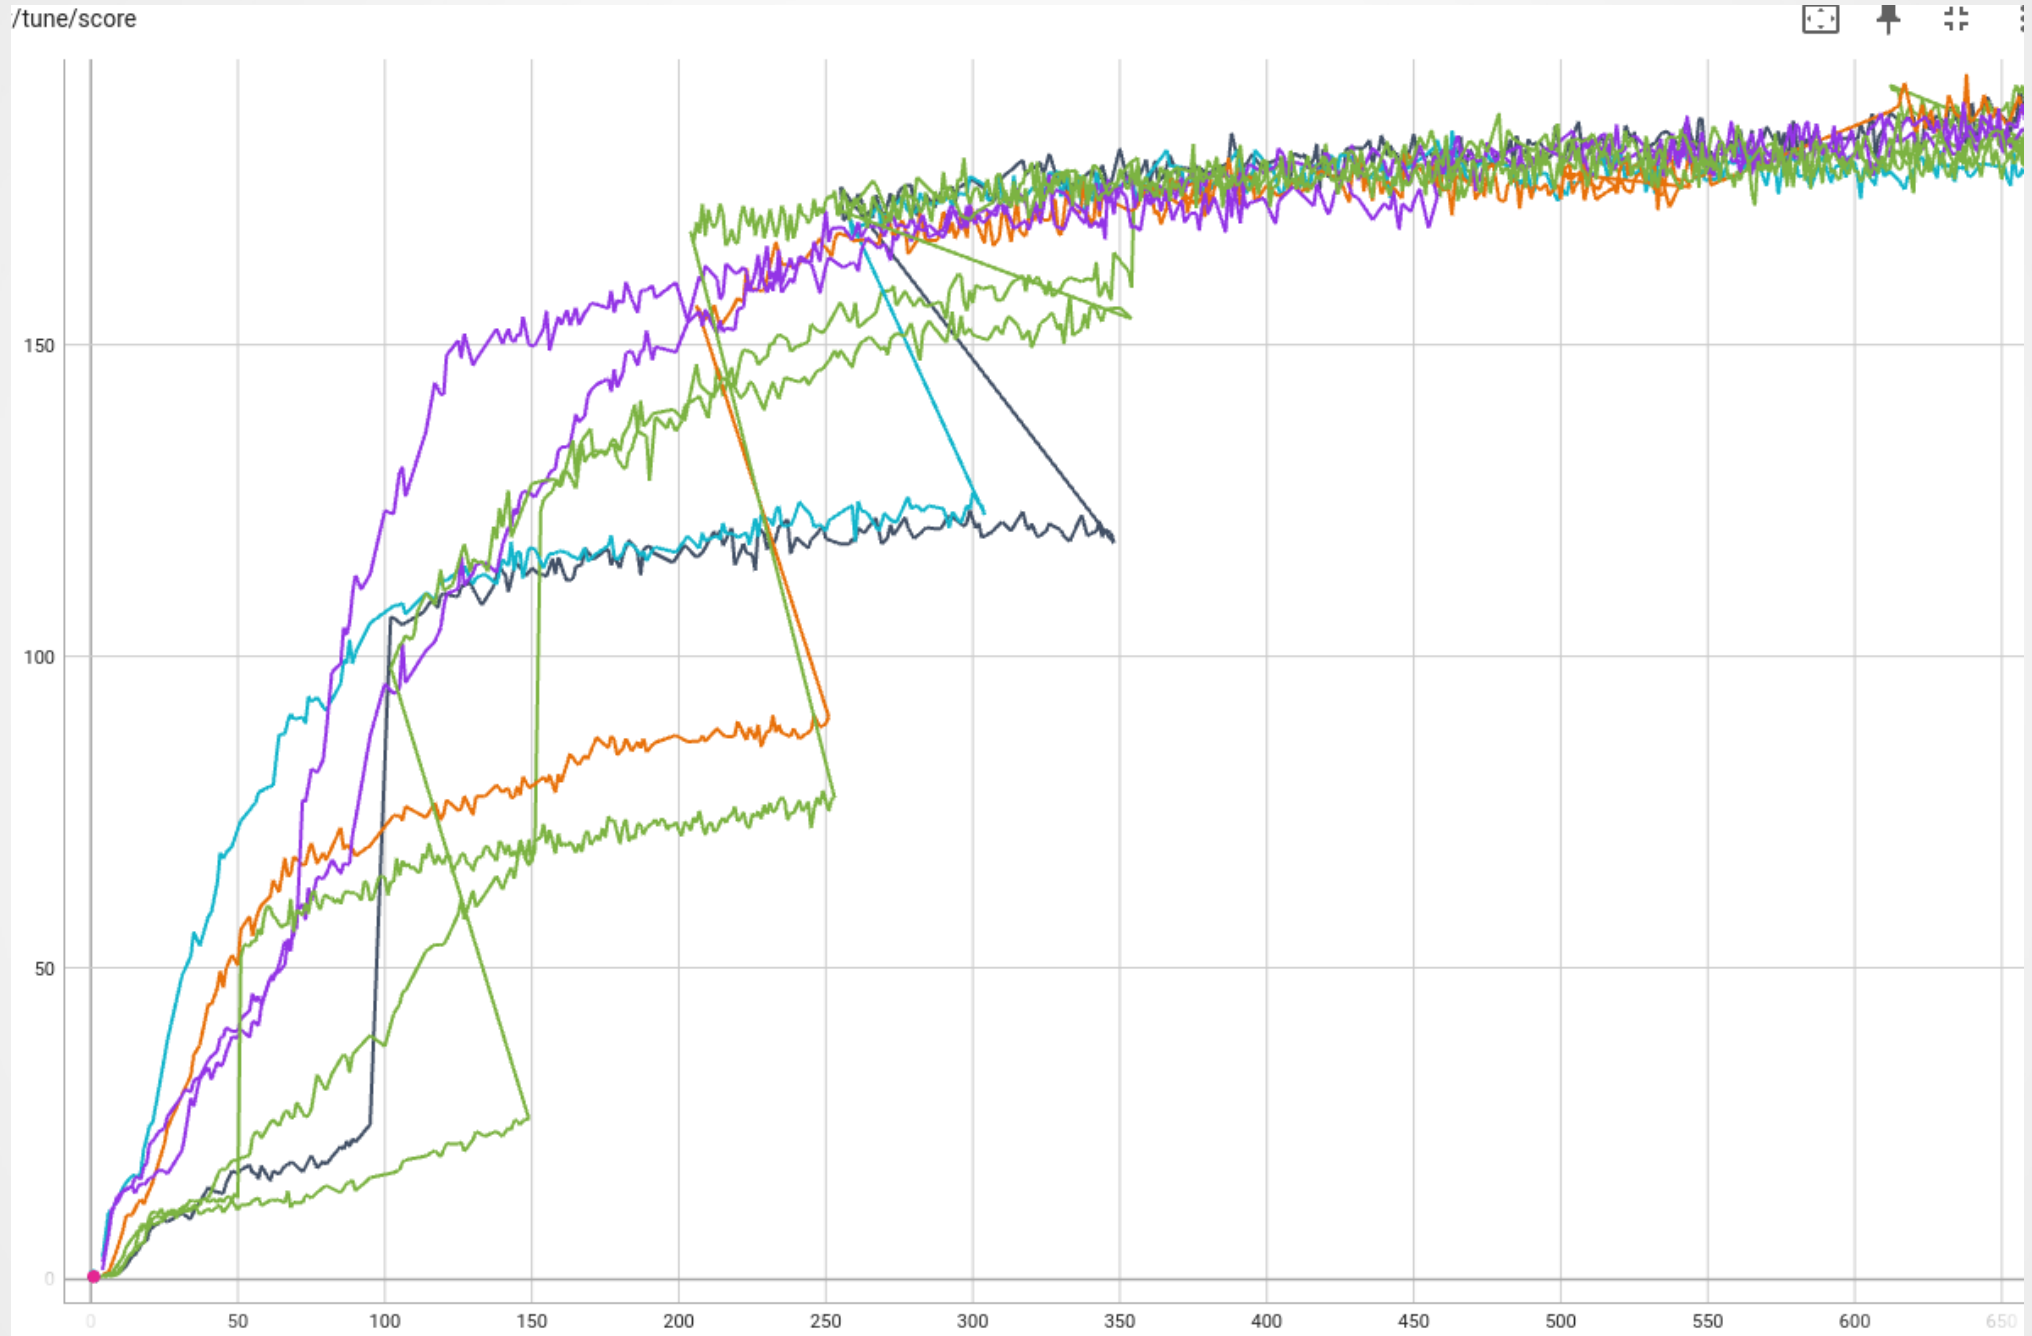
- “Population Based Training”
- **Replace** Bad Trials
  - Mutate Hyperparams
  - Copy Weights
- Popular for RL

...hyperparam sampling bias...

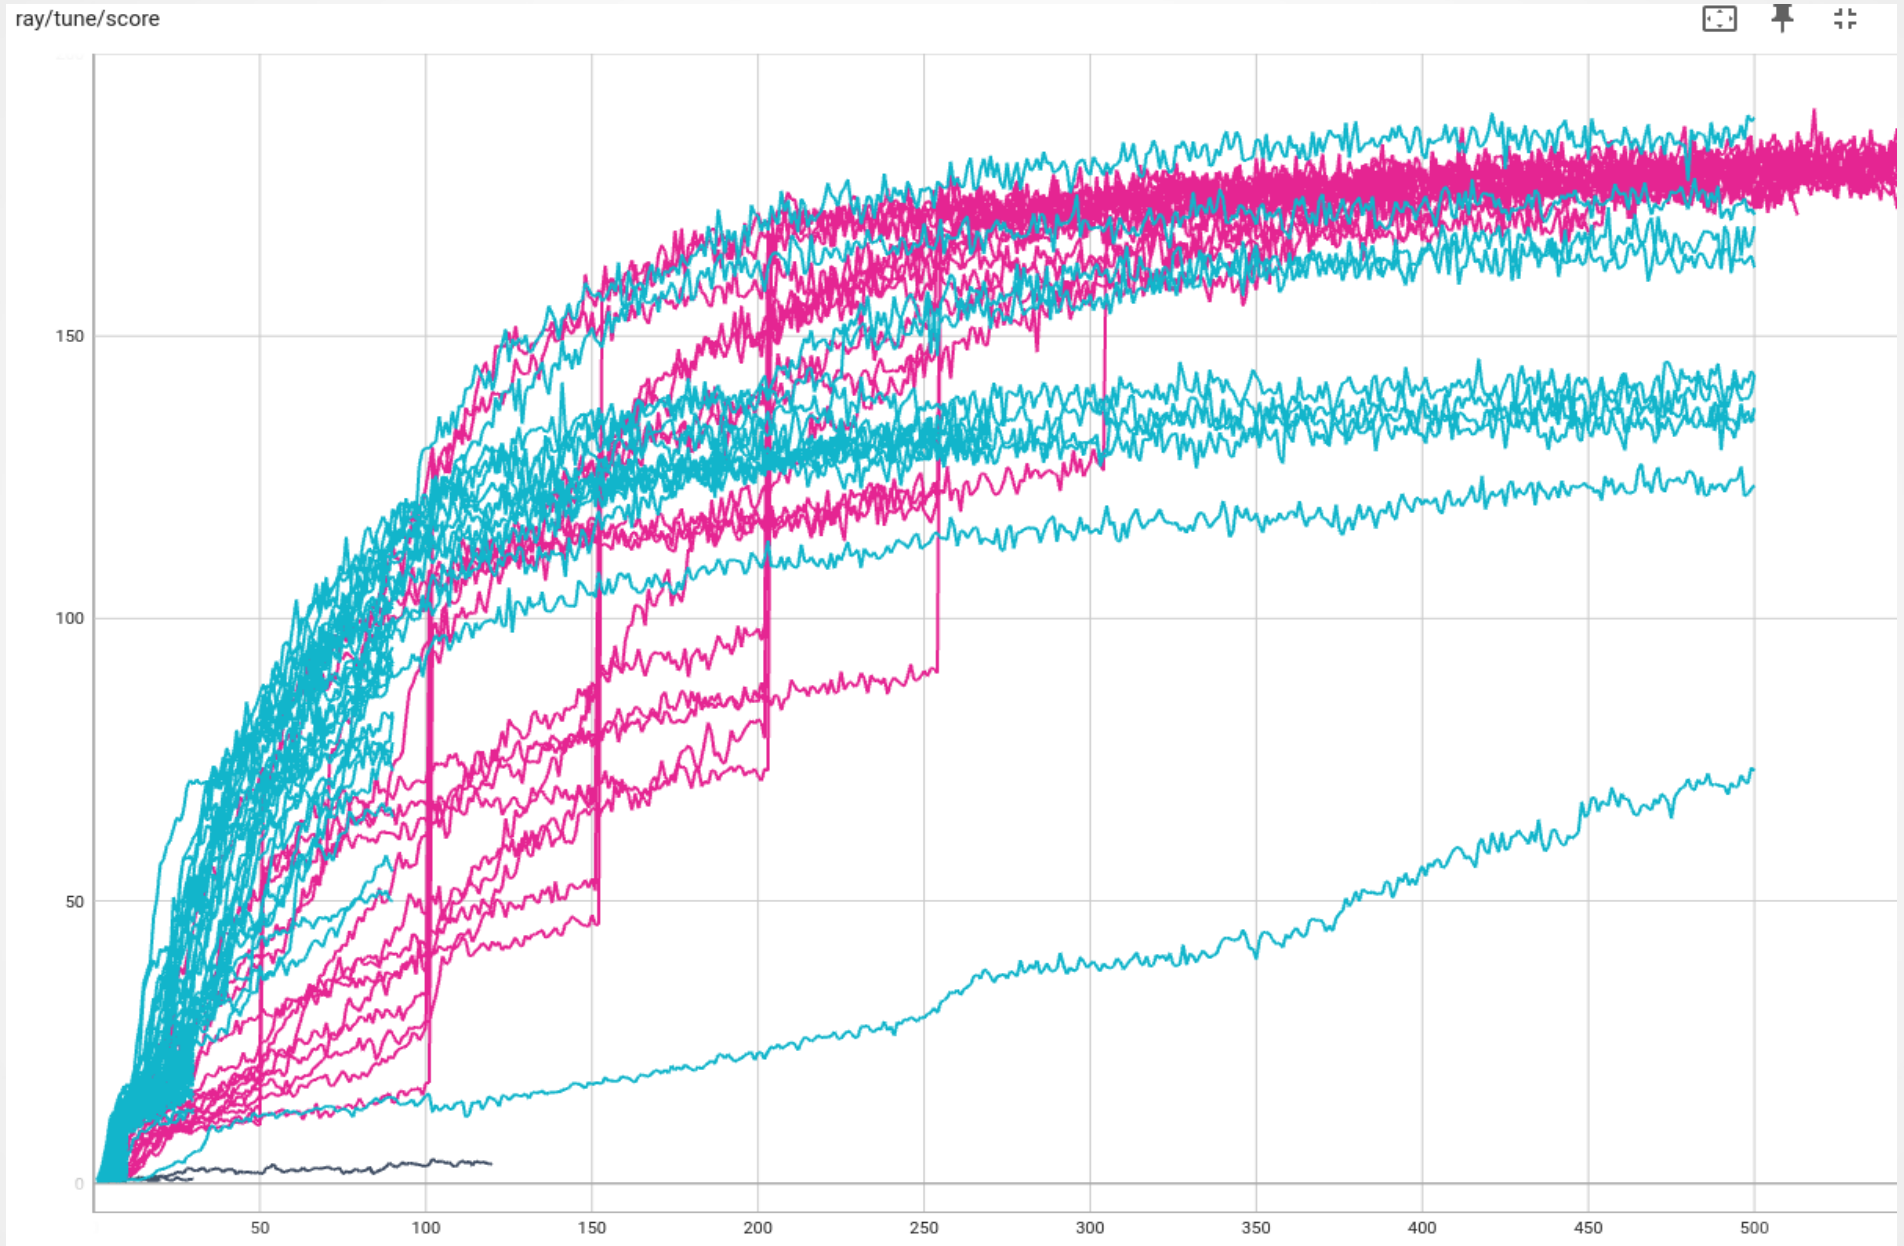
...ASHA as “restart”-strategy for CMA-ES? (vs IPOP/BIPOP)



# Population Based Training

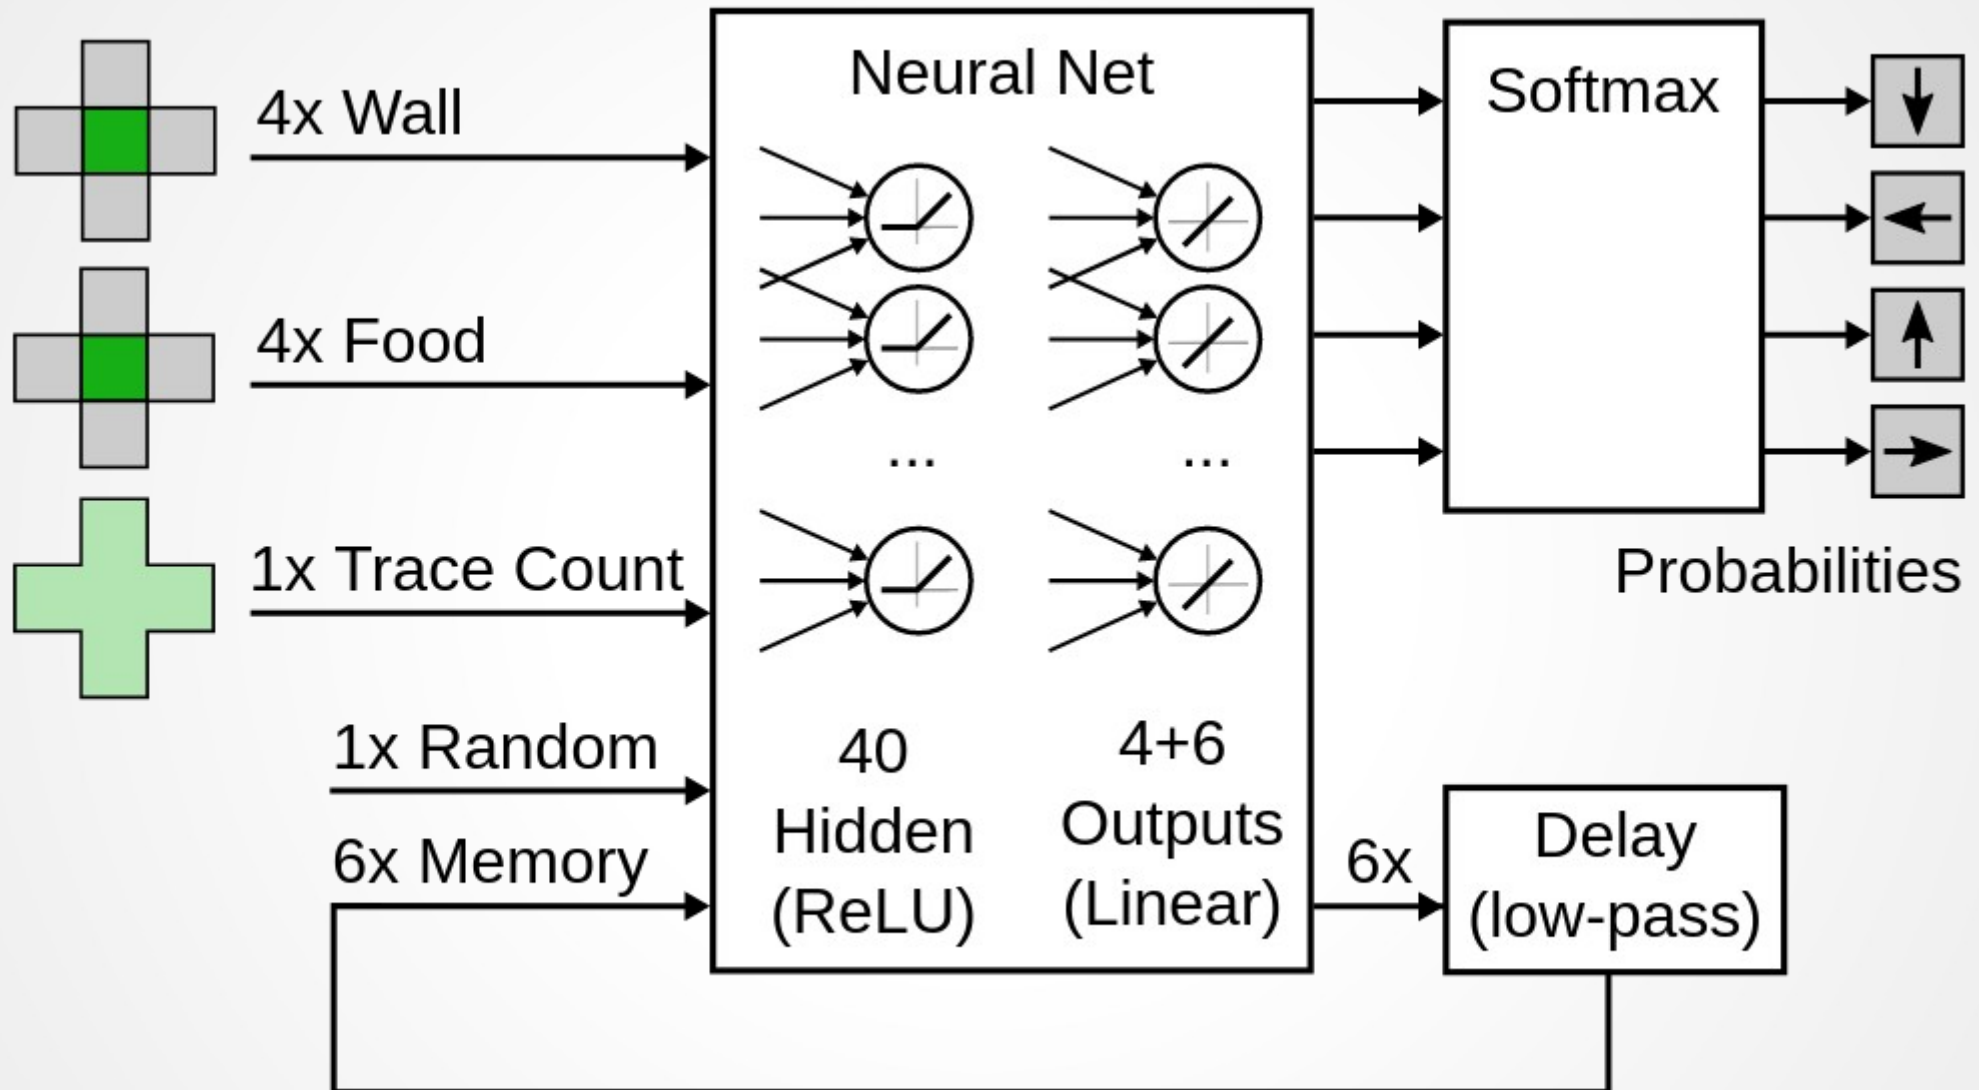


# PBT vs ASHA



End

# Policy



# Cluster Config

```
max_workers: 6
provider:
  type: aws
  region: eu-north-1 # Stockholm

available_node_types:
  ray.head.default:
    node_config:
      # required for s3 snapshots
      IamInstanceProfile:
        Arn: arn:aws:iam::909288963394:instance-profile/ray-head-v1
      ImageId: ami-0c03aae7ae21b81eb
      # head needs RAM (c6g.xlarge went OOM)
      InstanceType: m7g.xlarge
      resources: {"CPU": 0}

  ray.worker.default:
    node_config:
      # required for s3 snapshots
      IamInstanceProfile:
        Arn: arn:aws:iam::909288963394:instance-profile/ray-worker-v1
      ImageId: ami-0c03aae7ae21b81eb
      InstanceType: c7g.8xlarge # $0.24/h spot (2024-03) (64G RAM, 32vCPUs)
      InstanceMarketOptions:
        MarketType: spot
      resources: {"CPU": 34} # easier to saturate CPUs by over-subscribing
    min_workers: 0

rsync_exclude:
  - "target"
  - "node_modules"

head_setup_commands:
  - sudo fallocate -l 10G /swapfile
```

# Logging (PBT)

```
[PopulationBasedTraining] [Exploit] Cloning trial df83e_00007 (score = 1.572618) into trial df83e_00009 (score = 0.5469)
[PopulationBasedTraining] [Explore] Perturbed the hyperparameter config of trialdf83e_00009:
episodes_per_eval : 60 --- (* 1.2) --> 72
```

Trial BuildersTrainable\_df83e\_00002 started with configuration:

...

Trial BuildersTrainable_df83e_00002 config	
episodes_per_eval	51
popsize	120
rust/actions_scale	6.8
...	

```
(BuildersTrainable pid=15385, ip=172.31.40.141) Restored on 172.31.40.141 from checkpoint:
Checkpoint(filesystem=s3, path=maxy-ray-experiments/builders-pbt-7/BuildersTrainable_df83e_00000_0_2024-04-05_10-38-27/)
```

```
Trial status: 2 PAUSED | 8 RUNNING | 6 PENDING
Current time: 2024-04-06 10:02:04. Total running time: 32min 3s
Logical resource usage: 88.08000000000001/238 CPUs, 0/0 GPUs
Current best trial: 7a12c_00003 with score=57.9375 and params={'episodes_per_eval': 58, 'popsize': 120, 'rust': {'actio
```

Trial name	status	episodes_per_eval	iter	total time (s)	score	score_std
BuildersTrainable_7a12c_00002	RUNNING	182	38	1401.15	16.4453	4.66274
BuildersTrainable_7a12c_00003	RUNNING	58	39	1292.29	57.9375	11.9909
BuildersTrainable_7a12c_00004	RUNNING	112	38	1169.07	25.0781	6.02556
BuildersTrainable_7a12c_00005	RUNNING	226	20	720.775	10.3281	3.36497
BuildersTrainable_7a12c_00006	RUNNING	64	17	674.231	16.0117	4.96908
BuildersTrainable_7a12c_00007	RUNNING	329	19	635.388	8.14062	3.11808
BuildersTrainable_7a12c_00008	RUNNING	143	12	460.214	7.29297	2.94971
BuildersTrainable_7a12c_00009	RUNNING	297	2	77.3141	0.539062	0.355533
BuildersTrainable_7a12c_00000	PAUSED	259	50	1524.95	11.1211	3.84978
BuildersTrainable_7a12c_00001	PAUSED	122	50	1577.88	50.418	10.6566
BuildersTrainable_7a12c_00010	PENDING	80				
BuildersTrainable_7a12c_00011	PENDING	191				