Dumb Imitation

Is there a type of learning which is evolutionarily motivated, and which is not reinforcement learning?

Evolutionary motivation

This is a *theoretical* argument about the possible *evolutionary role* of a simple type of imitation learning.

I am *not* going to consider how imitation is cognitively achieved: it could be by sophisticated cognitive mechanisms, or something as simple as taking a rest at the same time of day, or travelling at the same speed.

Imitation can be adaptive for several reasons: these will *not* be considered

Amortization of exploration and trial-and-error

Adaptive to imitate those who have found the food, or acquired useful skills by trial and error.

Amortization of selection

Adaptive to imitate behaviour of those who have survived selection: behave like a survivor.

Imitation of parents as cognitive development

Transmission of culture from one generation to the next

Inter-generational amortization of trial-and-error learning.

Dumb imitation of peers

- An animal *partially* imitates the behaviour of a sample of its peers.
- It does *not* identify which behaviours are adaptive
- It does *not* identify which individuals are fitter: it imitates the behaviour of fit and unfit alike.

A hypothetical species: summerpodes

(A bit like the Megapodes …)

Four stage life cycle:

1. Egg-stage
2. Observation-stage
3. Learning-stage
4. Selection-stage

Summerpode life-cycle: egg stage

1. Egg-stage

   Eggs are laid in autumn and hatch in the spring.

   Parents die when eggs are laid, so next generation never meet previous generation, hence no cultural transmission. Each generation only meet each other.

   The **only** information transmitted from parents to children is **genetic**.
2. Observation-stage
3. Learning-stage
4. Selection-stage

Summerpodes: observation stage

1. Egg-stage
2. Observation-stage

   After hatching, summerpodes gather in groups of unrelated individuals.

   Each individual's behaviour is innate.

   The summerpode hatchlings observe each other's behaviours, without learning.

   There is no selection at this stage: all individuals survive.
3. Learning-stage
4. Selection-stage

Summerpodes: learning stage

1. Egg stage
2. Observation stage
3. Learning stage

   At the end of the observation stage, the summerpodes go to sleep; while asleep, learning occurs using their observations.

   Each summerpode adjusts its innate behaviour to be more similar to the the behaviour of its peers.

   When the summerpodes wake up, learning is complete, and their behaviour is now fixed for the rest of their life.
4. Selection stage

Summerpodes: selection stage

1. Egg stage
2. Observation stage
3. Learning stage
4. Selection stage

The summerpodes compete to feed, mate, lay eggs etc.

Reproductive success depends on the behaviour fixed in the learning stage, not on the original innate behaviour

Summerpodes die at the end of this stage.

Summerpodes: points to note

- The **only** learning is imitation of innate behaviour of peers in the same generation; there is no learning from experience
- During observation and learning, the summerpodes cannot distinguish which behaviours are adaptive, nor can they tell which individuals will be fit
- In this sense, summerpode learning is "dumb imitation of peers"

Summerpode genetics and evolution

Phenotype (learned behaviour) is abstracted as a vector of N real numbers.

Genotype is a vector of N real numbers.

Breeding: child genotype c is recombination of parental genotypes, plus Gaussian noise.

Selection: fitness decreases with distance of phenotype from ideal target phenotype t.

Summerpode learning

Each summerpode observes the innate behaviour of k randomly selected individuals.

"Observing innate behaviour" is abstracted as observing k genotypes $g_1 \dots g_k$

Each summerpode computes average of k genomes it observes, and adjusts its own innate behaviour towards this average.

$$p = (1-\lambda)g + \lambda \frac{g_1 + \cdots + g_k}{k}$$

A population of 2-dimensional summerpodes: blue dots represent genotypes. Yellow dot is one genome for which learning will be diagrammed.

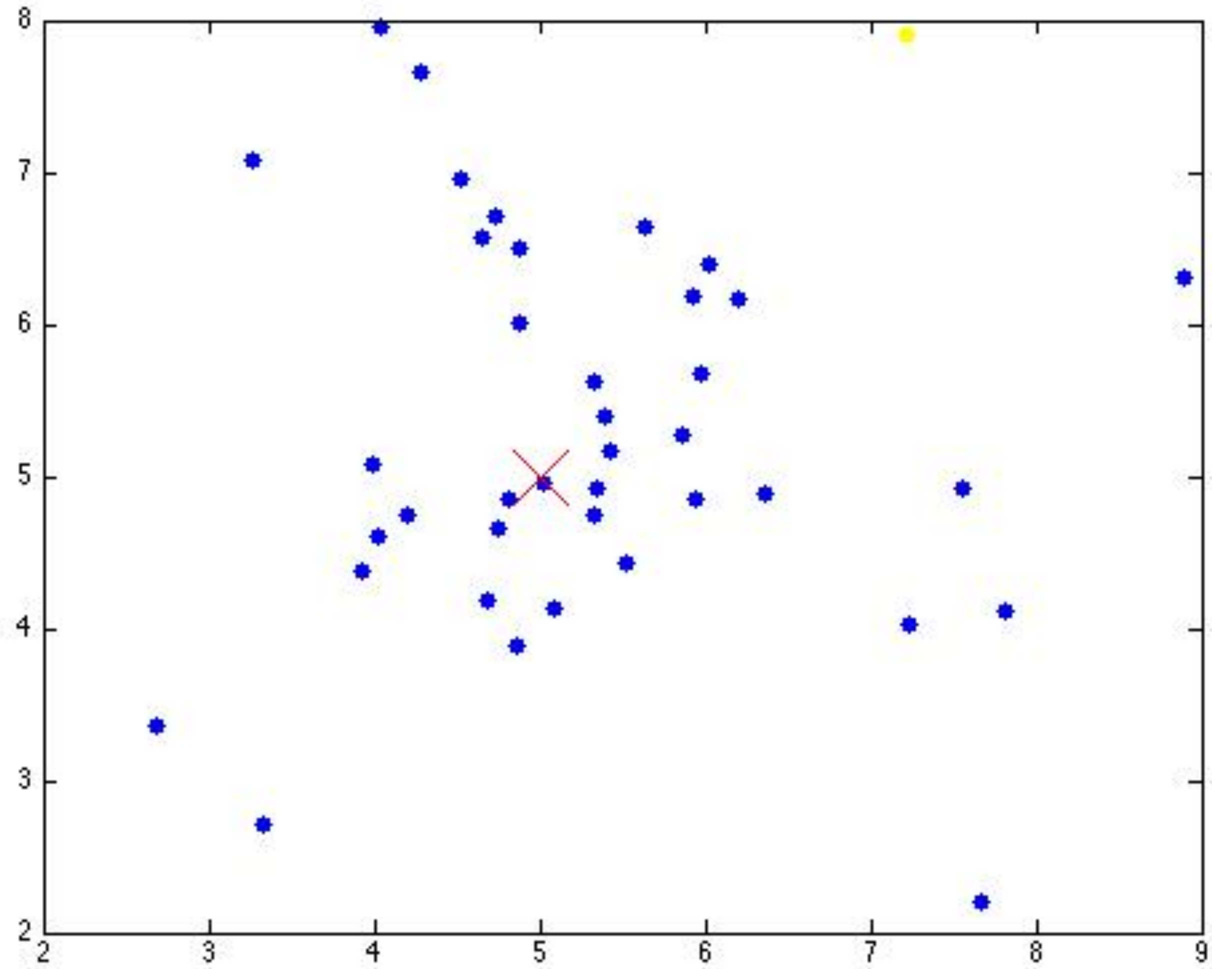

The target value (fittest possible phenotype) is shown by a magenta cross.

The yellow summerpode observes a sample of 10 others, and computes their average, as the cyan dot. Each summerpode in the sample of 10 is connected to the average with a magenta line.

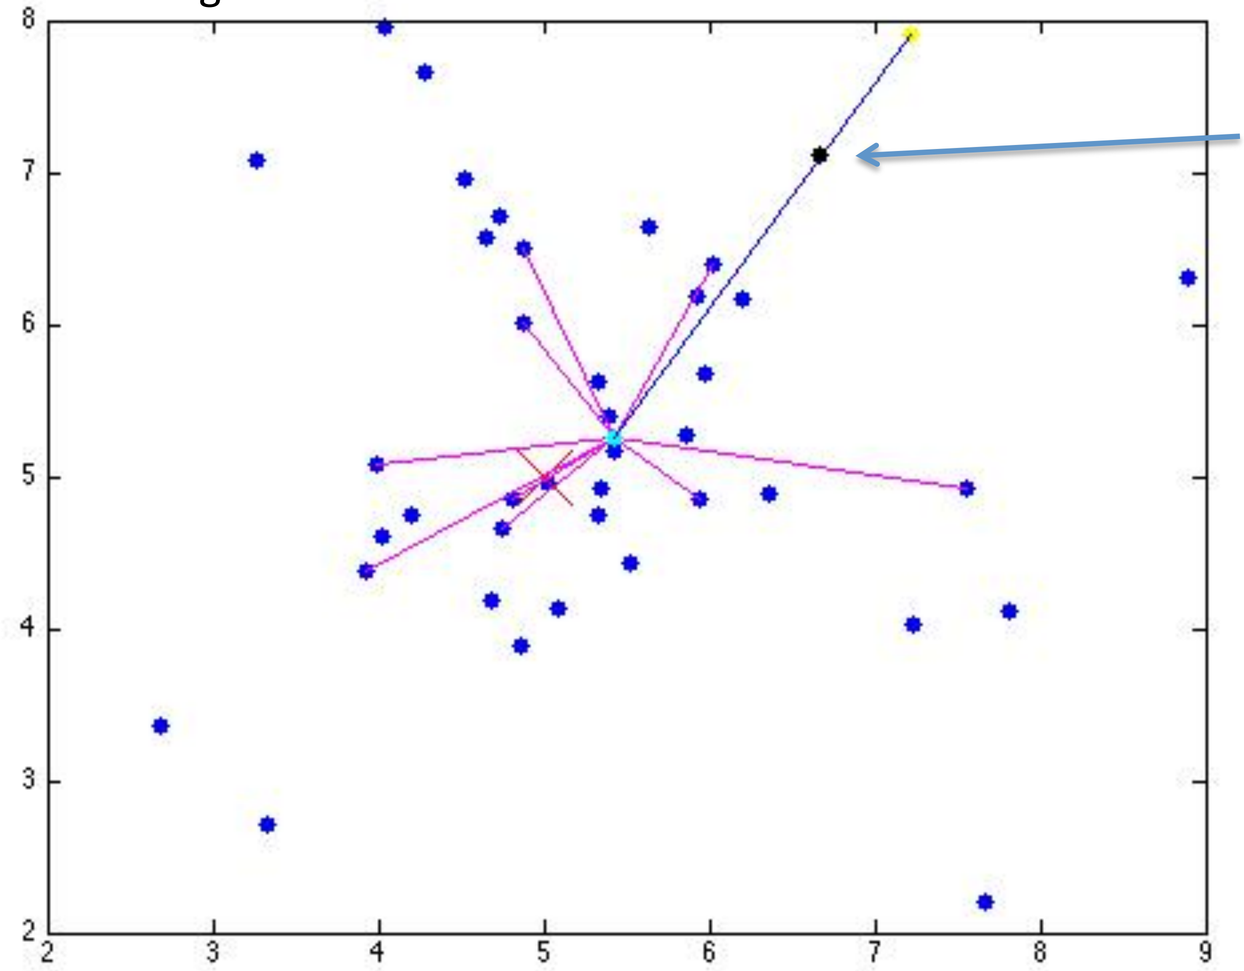

Learned Phenotype, a linear combination of the genotype and the average of the observations

The black dot (indicated by an arrow) on the line between the observer (yellow) and the mean of the observations (cyan) is the new learned phenotype for the yellow individual.

In this diagram, all the learned phenotypes are shown as black dots. Note that these are generally closer to the target than the genotypes.

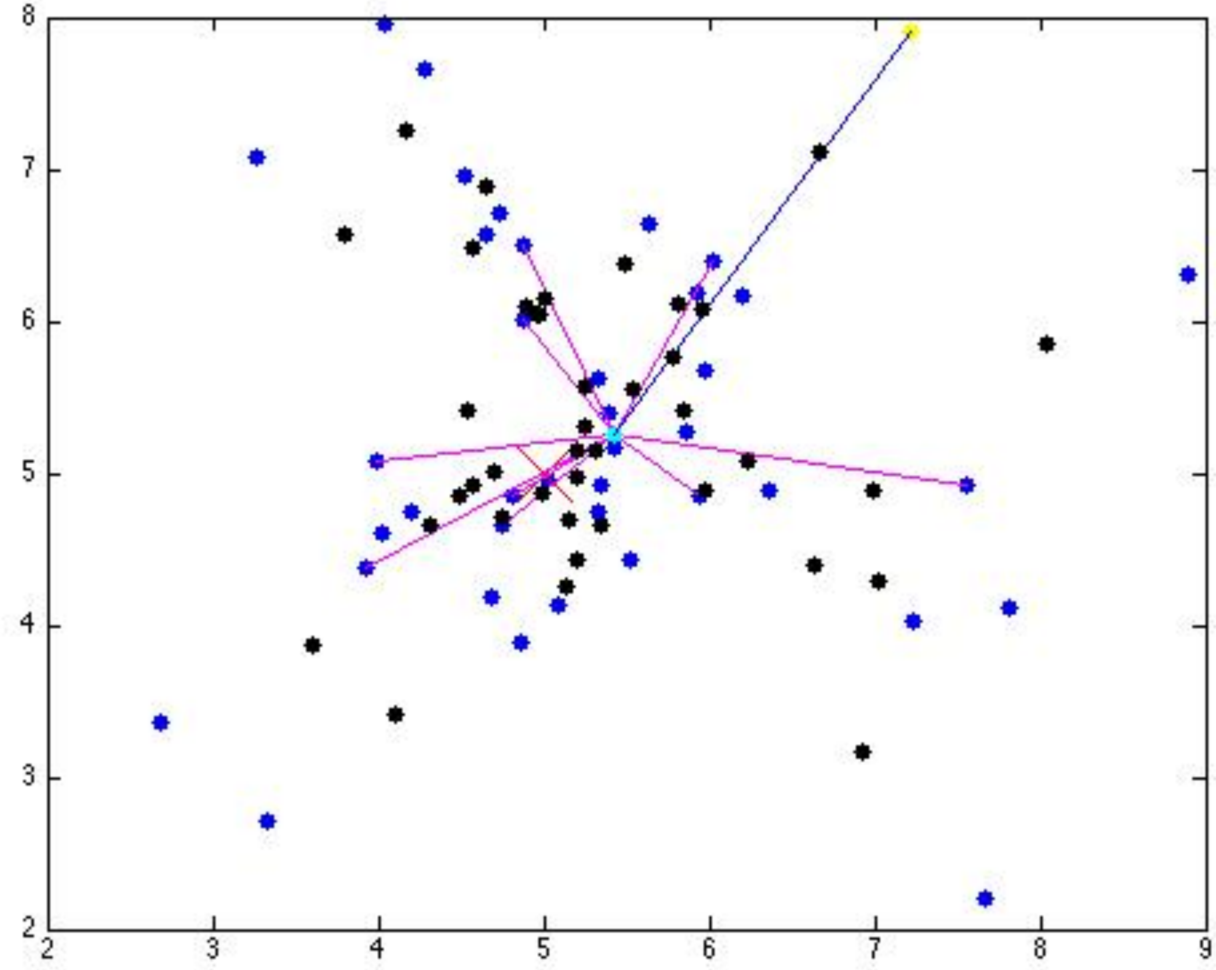

Summerpodes are now selected according to their learned phenotypes. The 50% of individuals with learned phenotypes closest to the target are selected.

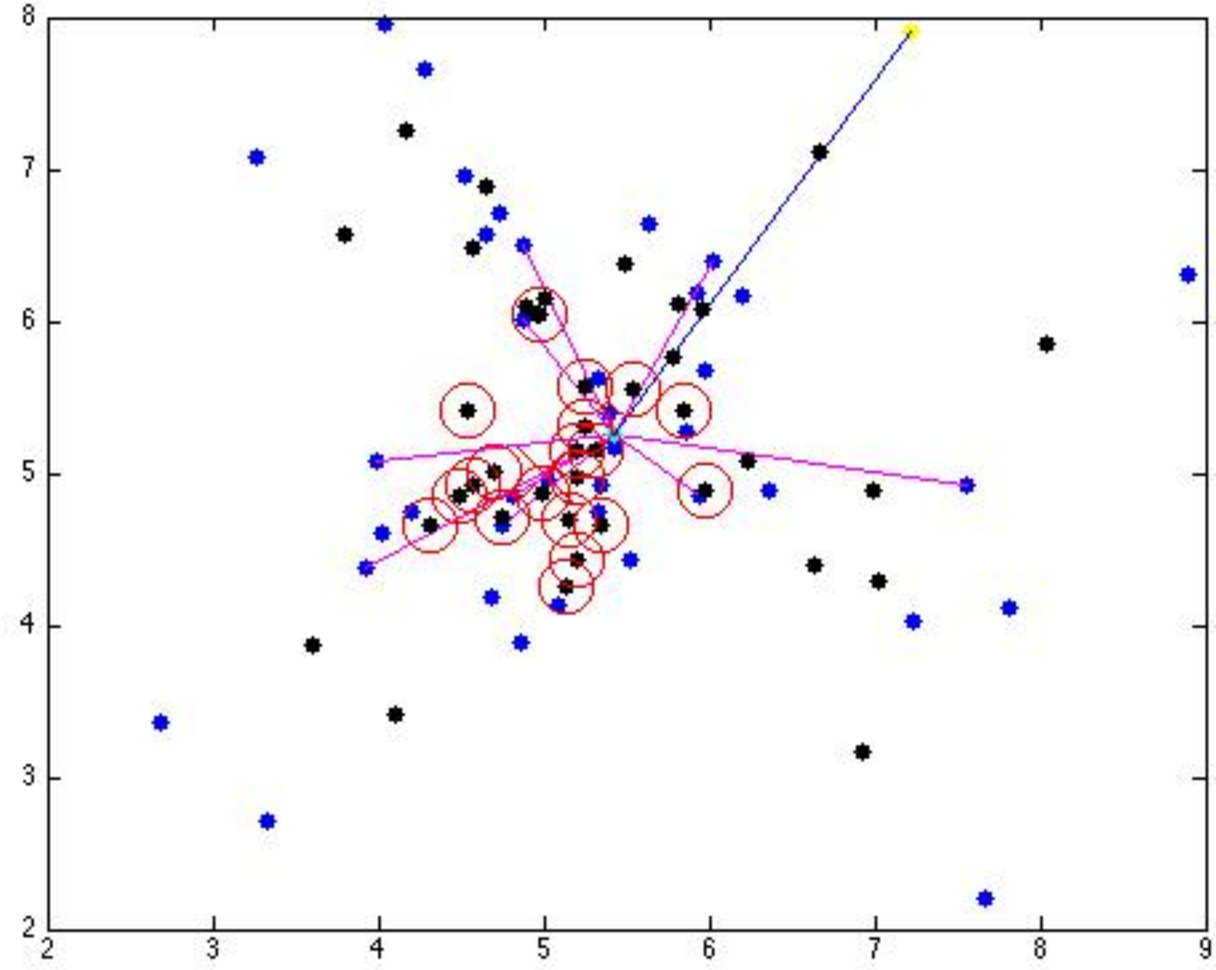

This diagram shows as red dots the genotypes corresponding to the selected phenotypes. It is these red genotypes that will be used to breed the next generation.

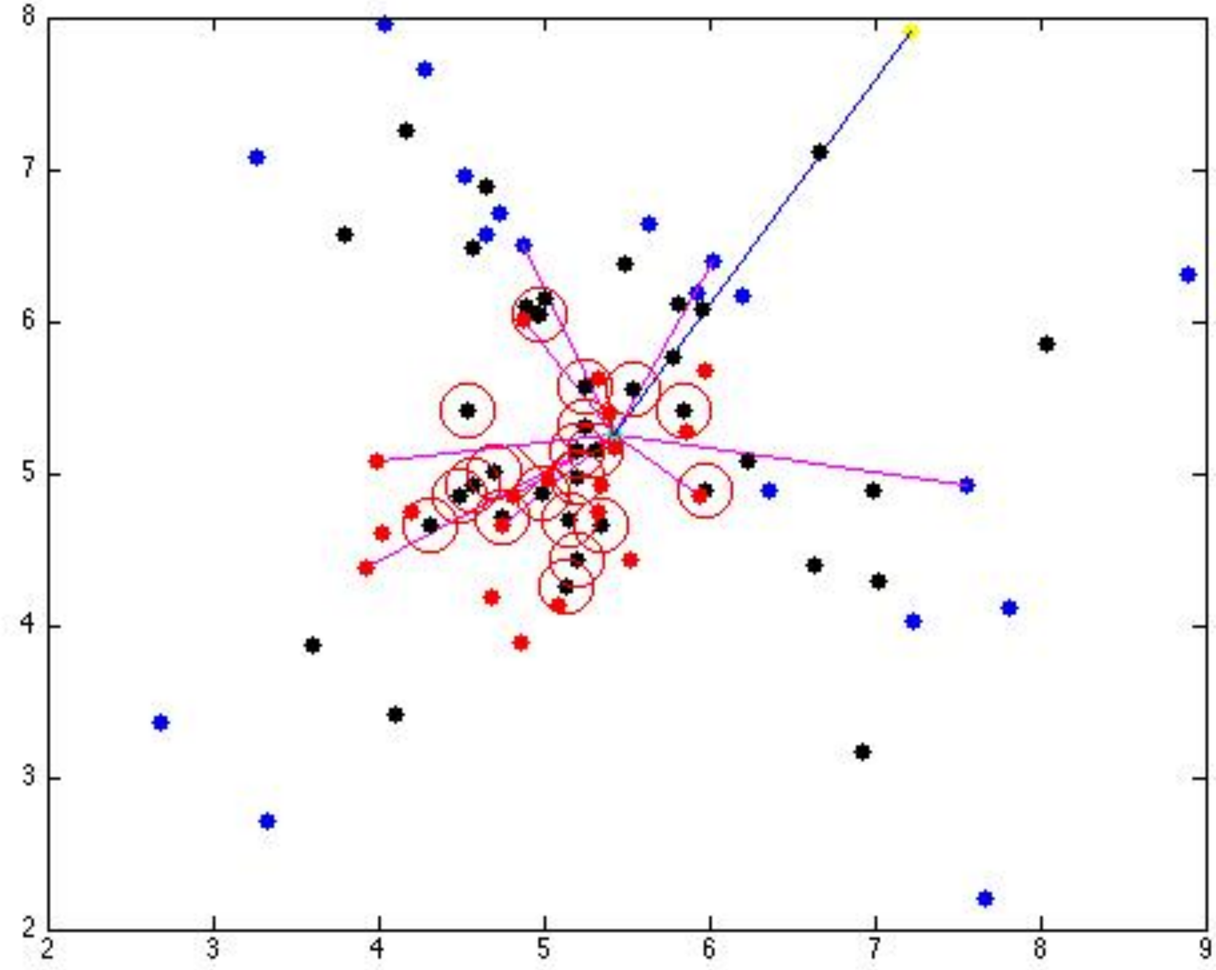

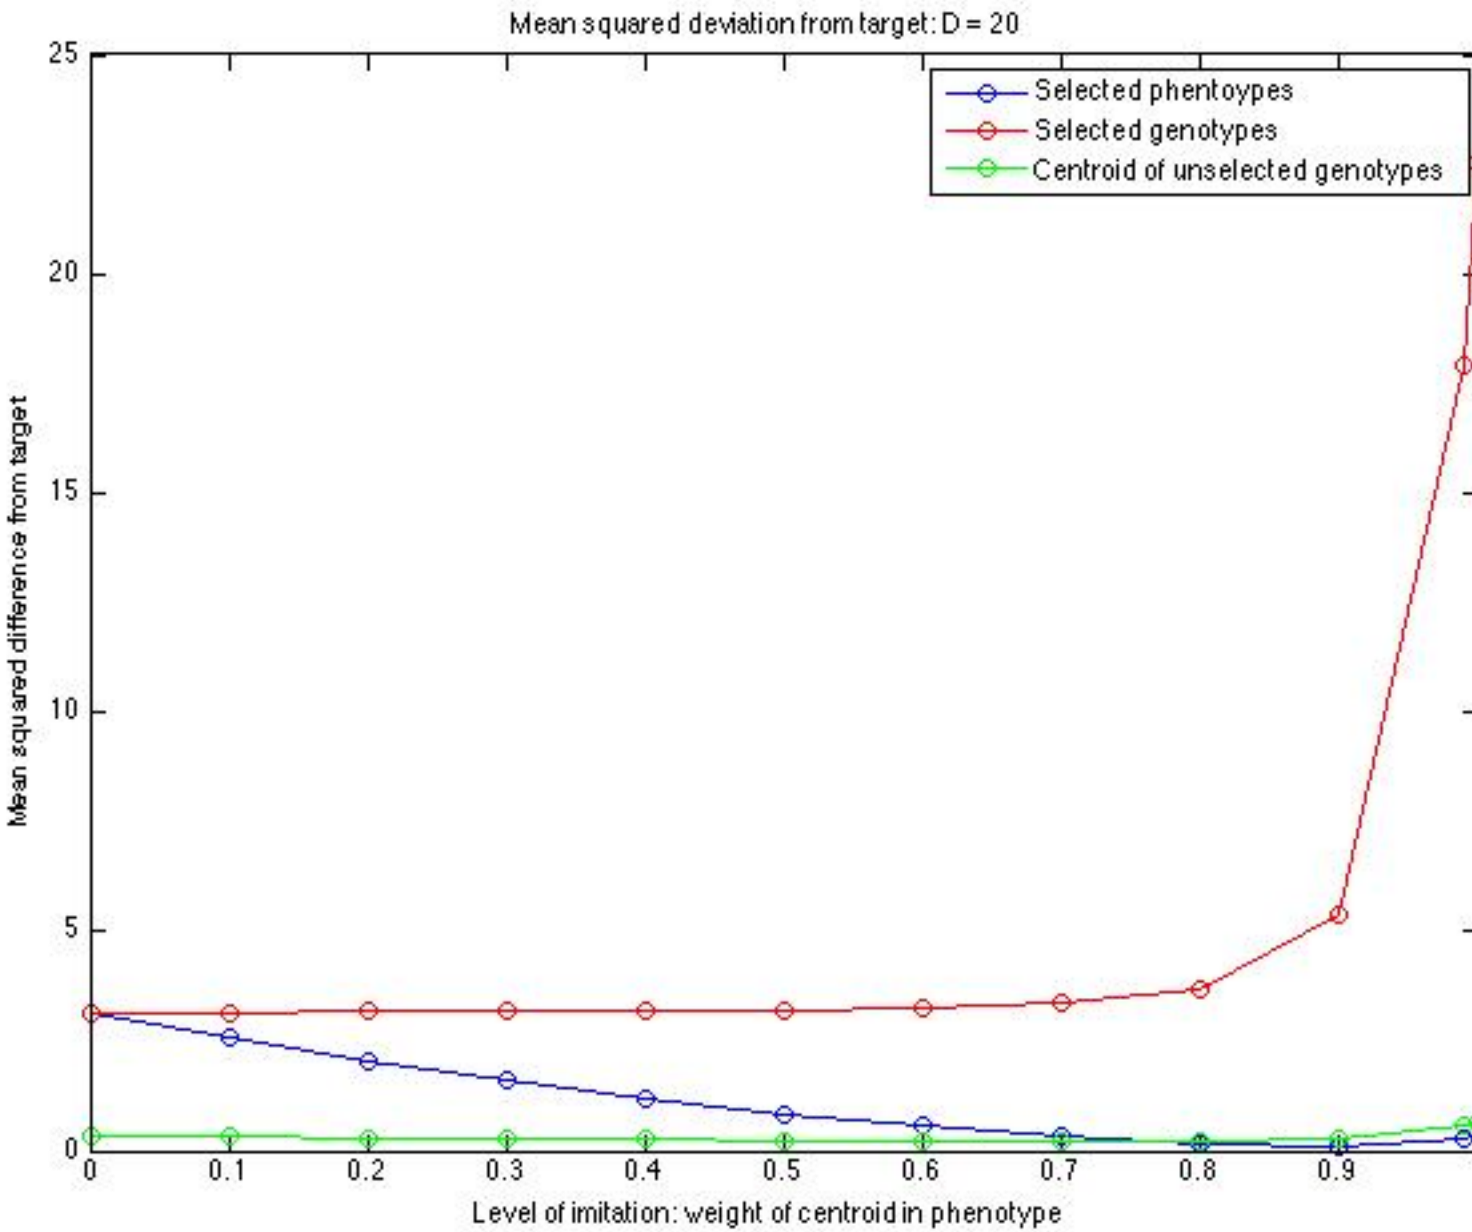

Mean squared deviation from target: D = 20
Selected phentoypes
Selected genotypes
Centroid of unselected genotypes
Mean squared difference from target
Level of imitation: weight of centroid in phenotype
25
20
15
10
5
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Remarks on the summerpode model

- Model has been greatly simplified to eliminate other forms of imitative learning, and use dumb imitation only.
- Dumb imitation might combine positively with other types of “smart” imitation.
- Simultaneous observation and learning would be a more realistic assumption, and would lead to a more complicated learning equilibrium, in which the effective value of k would be increased.
- Relatedness effects and non-uniform mixing during observation would reduce the effective value of k.

Population-level information flow

With dumb mutual imitation of innate behaviour, *less* selection gives *more precisely adapted* behaviour than with innate behaviour alone.

Yet the imitation is "dumb" in the sense that imitation takes no account of fitness.

So where is the extra information coming from ?!

Answer: **evolution is happening at the population level**.

The population *as a whole* contains more adaptive information than a single individual: imitation allows individuals to access this population-level information.

Conclusions on Dumb Imitation

1. Dumb imitation can enable an individual to access some types of genetic information across the population by copying innate behaviour
2. Equilibrium after learning can give more precisely adapted behaviour with less intense selection
3. Difficult to disentangle dumb imitation from smart imitation experimentally….
4. ….but dumb imitation could be ubiquitous in social species?

Reflection

Evolution is barely studied within machine learning, and yet evolution is a robust adaptive process – a kind of reinforcement learning – that has developed chemical, mechanical, anatomical, and behavioural solutions to the challenges of survival.

Two abstract models that relate evolution to learning have been briefly described.

There are surely many more to find...