University of Groningen

Modelling the evolution of learning

Kozielska, Magdalena; Weissing, Franz J.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2022

Modelling the evolution of learning

Magdalena Kozielska & Franjo Weissing

Groningen Institute for Evolutionary Life Sciences, University of Groningen, The Netherlands

m.a.kozielska@rug.nl

Introduction

Learning from experience leads to a change in the nervous system that manifests as altered behaviour. The ability to learn is an important adaptation, but how natural selection shapes learning is still not well understood. Evolution of learning is often studied using over-simplified analytical models or biologically unrealistic machine learning methods.

Here, **we present a novel way of modelling the evolution of learning using small neural networks and a biology-inspired learning mechanism**. Our learning mechanism is motivated by the role of dopamine in learning when expectations are not in line with reality. We use this model to answer the following research questions:

In which environments does learning evolve?   How efficient is evolved learning?   What is the effect of lifespan on the learning strategy?

Methods overview

We used individual-based simulations to study the evolution of simple, but capable of learning, neural networks experiencing variable environments.

Environment – Individuals live in variable and non-homogenous environments. During their lifetime they need to make multiple choices on where to forage.
**Task**: Choose a foraging patch that is of the highest quality among the available options. To judge the quality individuals use environmental cues.
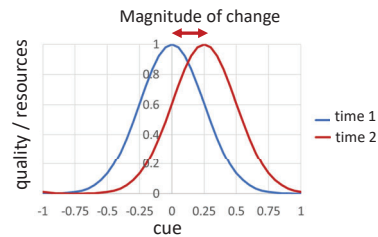
The quality of foraging patches is indicated by cues. The association between cues and patch quality remains the same within a lifetime but can change between generations.
Here: at time 1 (blue curve), cue values around 0 indicate highest patch quality, while cue values around 0.25 indicate highest patch quality at time 2 (red curve).

A heritable **neural network** is used to estimate the habitat quality for a given cue. Part of the network can be adjusted during individual lifetime in response to experience.
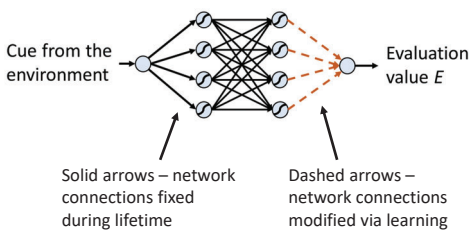
Learning occurs via a change in some network connections, making use of the "Delta Rule":

$$\Delta w = error * L * a$$

error: difference between real habitat quality and estimated habitat quality
L: **learning rate** (speed of learning)
a: activation of the preceding node

Solid arrows – network connections fixed during lifetime
Dashed arrows – network connections modified via learning

Individual lifetime consists of a learning period, followed by a foraging period.

Learning period: each time step: one learning opportunity using a random cue

Foraging period: each time step: evaluation of a subset of *k* options; choosing one with the highest evaluation value *E;* gaining resources equal to the real quality of the chosen environment

Reproduction – individuals with networks and learning strategies that lead to better choices (more resources gained during the whole foraging period) have more offspring.

Evolving parameters: network weights, duration of the learning period (**number of learning episodes, LE**) and **learning rate L**. The latter two are the focus of this poster.

## Results 1 – Example of evolutionary trajectories and learning ability

When the environment changes frequently, efficient learning evolves readily: both number of learning episodes and performance (ability to choose the best environment) increases and reaches steady state in short evolutionary time.

A relatively complex learning task can be efficiently solved by the evolution of a simple learning mechanism.
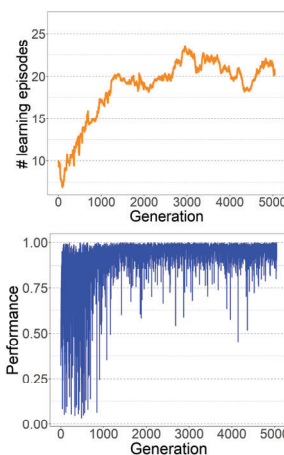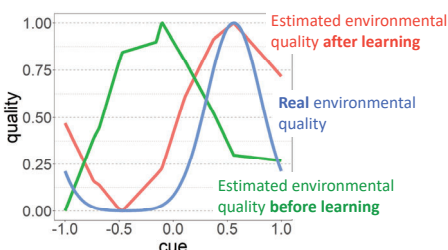

Results 2 – Effects of environmental change on learning

For a **lifespan of 500** time steps, the figures show the evolutionary outcome for three frequencies of environmental change (separate panels) and three magnitudes of change (x-axis). Shown are 20 replicate simulations per parameter combination and each point represents the population average of one replicate.
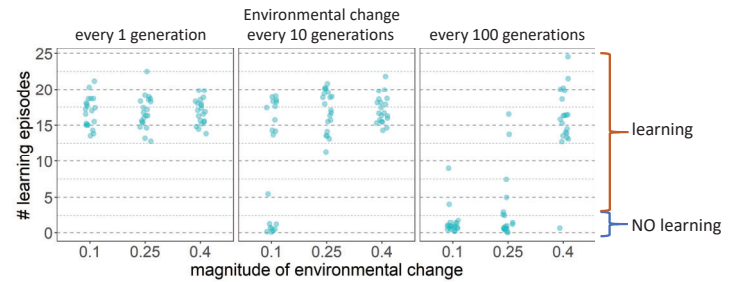

Learning readily evolves if environmental change is sufficiently large and/or frequent. When learning evolves, the length of the learning period is largely independent of the frequency and magnitude of environmental change.

Simulations in which learning (LE) can evolve

Simulations without learning, LE fixed to 0
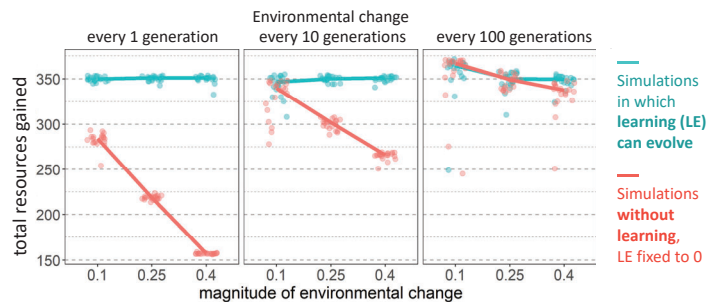
Even though learning is costly (time spent learning shortens the foraging period) the benefits of learning (better decisions in foraging period) outweigh the costs when environmental change is sufficiently large and/or frequent.

Results 3 – Effects of lifespan on the evolution of learning

Lifespan: 50 time steps
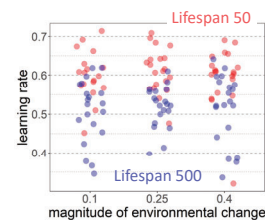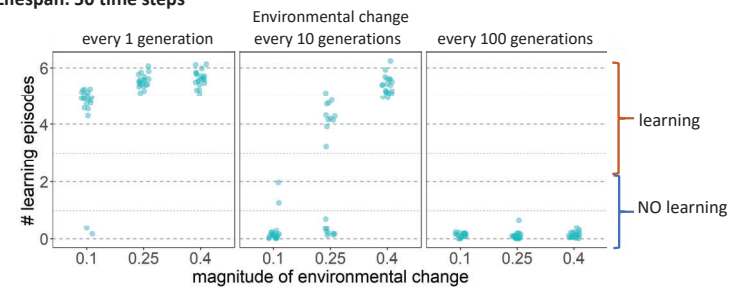
For shorter lifespans learning is less likely to evolve.

When learning evolves, the learning period is shorter in absolute terms, but it corresponds to a larger fraction of lifetime.

In case of a shorter lifespan, a shorter learning period is partly compensated by a larger learning rate.

General conclusions:

Efficient learning readily evolves in a model based on simple learning mechanism.
The degree of environmental variability affects the likelihood that learning evolves but not the evolved learning strategy.
Lifespan has a substantial effect on the evolution of learning and the evolved learning strategy.