

Adaptive Growth Processes: A Model Inspired by Pask's Ear

Nathaniel Virgo¹ and Inman Harvey

University of Sussex, UK
¹nathanielvirgo@gmail.com

Abstract

We introduce the notion of an “adaptive growth process” in order to explain an experimental result from the 1950s in which a complex mechanism capable of distinguishing between two sounds emerges from a homogeneous chemical solution. We present a very simple computational model which exhibits an adaptive growth process. Adaptive growth processes could have practical applications in adaptive control systems and may also play a role in biological development.

keywords: self-organisation, complex systems, reinforcement learning, Gordon Pask

Introduction

In the late 1950s Pask (1958; 1960; 1961) was able to construct a device whereby a complex functional structure would emerge within a physical medium (a solution of ferrous sulphate), simply by increasing its supply of electrical current according to its performance at a given task. The task Pask set the device was to react to sound. The structure it produced resembled an ear, with an array of resonating metal threads which was able to distinguish between two different sounds. He was also able to get the same device to respond to changes in a magnetic field.

Unfortunately Pask did not record all the details of his setup and the experiment has never been repeated. Nothing quite like it has been achieved before or since. However, if the result can be repeated and generalised its potential would be enormous. It would mean that complex structures with functional components could be grown *in situ* without the need for an evolutionary population. For instance, one could imagine Pask's ear being used as a sensor in an adaptive robot. If the robot found itself in a situation where magnetic fields were relevant it could find itself adapting its ear to respond to them. If Pask's result is sufficiently general one could go further and imagine a neural controller that is able to grow adaptively, creating new neural structures in response to novel challenges. It is clear from Pask (1960) that Pask thought along similar lines.

In this paper we elaborate on the possibilities that this simple experiment opens up. We present our ideas on the

general principles behind the result, defining a general notion of an *adaptive growth process*. We back this up with a simplified computational model of an adaptive growth process which, although it does not perform as impressive a task as Pask's ear, works in what we believe is a similar way.

Our system resembles a model of ants leaving pheromone trails to reach a target, except that we do not selectively reinforce the trail of a successful ant. Instead the reward for reaching the target is applied to the system as a whole, in the form of an increase in the rate at which ants enter the system. This stems from a difference in focus between our model and typical ant or swarm based models. Our model is intended to illustrate a plausible mechanism behind the success of Pask's experiment, in which individual threads were not reinforced and in which the solution consisted of a network of many interacting threads.

We also explore the possible biological implications of this kind of emergent adaptive structure.

Pask's experiments

The electrochemical experiment in which the ear was produced is mentioned in Pask (1958; 1960; 1961) but is always presented as an example to back up a philosophical point rather than as an experimental result in itself and it is difficult to decipher the exact experimental conditions which were used. There is also the eyewitness account of Stafford Beer described in Bird and Di Paolo (2008), although it was given many years after the event and does not give a complete description. Further descriptions of Pask's experiment and the history behind it can also be found in Bird and Di Paolo (2008) and Carriani (2007). There is a photograph of the resulting “ear” structure in Pask (1958), which is reproduced in Carriani (2007).

Although the details are obscure the basic idea is clear. Pask was experimenting with passing an electric current through a solution of ferrous sulphate. This causes a thin metal wire to be deposited along the path of maximum current. Since these threads have a very low resistance they affect the electric field surrounding them and thus the growth of further threads. Such systems were a popular way to

model neural growth at the time.

By varying the current flow through an array of electrodes Pask was able to affect the growth of the threads. For instance, when activating one negative and two positive electrodes a wire forms that starts at the negative electrode and branches toward the positive ones. If one of the positive electrodes is switched off the wire moves so that both branches point towards the remaining positive electrode, with the branching point remaining stable. If part of the wire is then removed the gap gradually moves toward the positive electrode, with the wire ahead of the gap being dissolved by the acidity of the solution but the wire behind growing back in its place. The branching pattern is reproduced by the new wire. Details of this can be found in Pask (1960).

Pask then took his electrochemical system and subjected it to sound. At this point the details become less clear but the system was rewarded in some way by an increase in available current whenever it responded to the sound, presumably by forming connections between a particular set of the electrodes. After about half a day a structure was formed which was able to perform this task, and Pask then went on to train it to distinguish between two tones, one about an octave above the other.

The interesting thing is that this structure is fairly complex, with functionally differentiated parts that are not specified by the experimenter. In Pask's words, "the ear, by the way, looks rather like an ear. It is a gap in the thread structure in which you have fibrils which resonate with the excitation frequency" (Pask (1960)). This solution is truly system-level. It is not determined simply by the fittest individual thread but by a combination of many threads playing several different roles. Some act as vibrating fibrils while others are part of the supporting structure. Presumably the fibrils can become further specialised to vibrate when exposed to sound in a particular frequency range. This spontaneous division of labour is not a feature of most learning algorithms.

Pask's Ear as a Dissipative Structure

The system of metallic threads that forms in Pask's ear is a dissipative structure. That is to say, it is a kind of structure that exists as the result of an externally imposed flow of energy through a system. The threads can only form and persist when an electrical current is passed through the medium, and if the current stops the acidity of the ferrous sulphate solution will gradually dissolve them. The structure maintains its form through a balance of creative and destructive processes.

This seems to us to lie at the heart of its operation. The system is subject to continual fluctuations, in which filaments are lengthened or shortened, or new ones grown and old ones destroyed. But the threads are in competition with each other: a given amount of electrical current can only support a certain amount of filament, so any new structure

which increases the amount of metal thread can only become stable and be maintained if it causes a sufficient increase in the current flow.

Adaptive Growth Processes

We use the term *adaptive growth process* to refer to a system which operates in this way, where structures compete for some resource and those that contribute towards maintaining an increased supply of that resource tend to be more stable over time.

This definition may be refined at a later date but some general preconditions for an adaptive growth process are that there is a substrate and a resource whose availability depends in some way on the state of the system. The substrate is such that without any inflow of the resource it will decay to a homogeneous state, but that an inflow of resource enables structures to persist in the system. The dynamics of the system must be such that these structures compete for the resource and are continually subject to fluctuations, resulting in structures that contribute to an increased resource flow becoming more likely to persist over time than those that do not.

Of course, this does not happen in all circumstances in all possible substrates. This research is a first step towards identifying the specific requirements for an adaptive growth process to take place. Some insights gained from our model can be found in the discussion section.

One interesting feature of adaptive growth processes is that they can in some circumstances become more complex: Pask's ear presumably starts with a fairly simple network of threads and ends up with a relatively complex arrangement of resonating fibres. This increase in complexity occurs when an increase in resource flow can be achieved by a more complex structure. Since an increase in complexity will usually require an increase in resource use (to maintain a larger amount of metal thread, for instance) the structure will not generally become more complex than it needs to be.

Relationship to Reinforcement Learning and the Credit Assignment Problem

Reinforcement learning is a general term for a learning algorithm which can learn a task (usually a classification task) by being given a "reward signal" whenever it behaves in the desired way. Pask's ear can be seen as an example of this, with the supply of electric current acting as the reward signal. In general the supply of resource to an adaptive growth process acts as a reward signal.

From a reinforcement learning point of view, an interesting feature of Pask's ear and adaptive growth processes in general is that they avoid the so-called Credit Assignment problem. This is simply the problem of deciding which part of the system to reward for a successful behaviour. Since behaviour results from the dynamics of the system as a whole it can be hard or impossible to know which parts contribute

to a given behaviour and which detract from it. Pask's ear solves this problem in a very simple way: the reward is applied to the whole system in the form of an increase in resource availability. Since the system's parts are in competition with each other it is only the ones which contribute to this increased resource availability that will remain stable in the long run.

Relationship to Evolution by Natural Selection

Our proposed mechanism for adaptive growth bares some resemblance to the process of natural selection. In both cases a system is subject to small random variations which are more likely to persist if they increase a certain target property of the system (its fitness to its environment in the case of natural selection, or the rate of resource input from the environment in the case of an adaptive growth process). In both cases it is the behaviour of the system as a whole, rather than its individual components, that determines which variations are selected.

The primary difference is that in natural selection the selection takes place between a number of similar systems, whereas in an adaptive growth process there is only one system and the selection occurs over time. Or, if one prefers to think of a system like Pask's ear as being a population of threads then the selection takes place according to a population-level property rather than to individual fitness; but again there is only one population.

Both processes could be simultaneously relevant in living systems. Adaptive growth processes within individual organisms could be honed by natural selection operating on a longer time scale, for instance.

Implications for Biological Development

All the structures that occur within living organisms are also maintained by a flow of energy (ultimately provided by the organism's food) and will decay if that energy flow ceases. Perhaps some of the complexity of biological structures is formed and maintained by what we have termed adaptive growth processes. If this is the case then research into these principles could vastly increase our understanding of biological growth and development processes. Pask saw this potential, writing in Pask (1960) that "the natural history of this network [of metallic threads] presents an over-all appearance akin to that of a developing embryo or that of certain ecological systems."

It seems quite possible to us that nature would take advantage of this "design for free" whenever possible. For instance, an organism's genes would not need to specify every component of an ear, but simply to arrange for circumstances in which an ear will emerge in response to stimuli from the environment and a modulated supply of energy. Of course there will also be a strong element of genetic design in nature, but an element of adaptive growth could perhaps explain why organs can atrophy or fail to develop properly

when not used. Conceivably it could also be a factor in the enormous plasticity of the brain.

One could also imagine adaptive growth processes occurring at a larger scale in the development of an ecosystem. The idea that ecosystems act so as to increase or maximise the flow of resources through them (or some related quantity) is an old one in ecology, dating back to Lotka (1922) but has fallen out of favour in recent decades due to the lack of a convincing explanation. Perhaps the study of adaptive growth processes could help to illuminate this issue.

An Adaptive Growth Process in a Model

Rather than try to simulate the complex physics of electrochemical deposition we have developed a very simple computational model in which we have tried to capture the conditions required for adaptive growth to take place: there is a balance between creative and destructive processes and structures compete for a resource whose supply is globally changed in response to the system's performance at a task.

In Pask's ear metallic threads are deposited by a flow of electrons between electrodes. We have abstracted this to something which resembles a system of pheromone trails laid down by ant-like entities which move across a grid. We shall use the metaphor of ants and pheromones rather than electrons and filaments to describe our model in order to emphasise that it is not meant to be a physical simulation of Pask's electrochemical experiment. However, we consider the pheromone trails and the rate of arrival of ants to be analogous to the metallic threads in Pask's ear and to the electric current respectively.

It is important to note that, unlike most ant-based optimisation algorithms we do not selectively reinforce the trail left by a successful ant. Instead the 'reward' is applied to the system as a whole in the form of an increase in the rate at which ants arrive.

Specification of the Model

The system is divided into a grid 50 cells wide by 500 high. Each cell contains a floating point number representing the amount of pheromone present. These are initialised to zero.

The following two processes are then repeated a large number of times: first an ant enters at the top of the grid and moves towards the bottom, tending to follow any pheromone trails that are present according to the rules given below, and laying down a pheromone trail of its own. This is followed by a period of time during which all pheromone trails decay. This period of time, and hence the amount of pheromone decay that takes place, is adjusted according to a simple reward function. Note that the two time scales are separated: the ants are assumed to travel instantaneously as far as the pheromone decay time scale is concerned.

Each ant enters the system at a uniformly random position on the top row, and moves toward the bottom one row at a time according to the following algorithm:

1. Look at the amount of pheromone in the cell directly below and the cells to either side of it, three cells in total (the edges wrap around).
2. Add the value $\delta = 0.1$ to each of these three numbers (this is to prevent very weak trails from having a strong effect) and normalise them to get a probability distribution.
3. Move to one of the three cells below according to the computed probability distribution
4. Add the value 1.0 to the amount of pheromone in the previously occupied cell

After an ant has reached the bottom there is a period of simulated time in which the pheromone in each cell decays exponentially. This amounts to multiplying each cell in the grid by the same number between zero and one. The length of this time, and thus the amount of decay, depends on the reward function, which is described below. In this way the rate of build up of trails can be controlled by modulating the time that elapses between each ant entering the system.

In addition to this, before each ant enters the system the value $d = 0.01$ is subtracted from the amount of pheromone in each cell (we set the value to zero if it becomes negative). This has a proportionally greater effect on weaker trails, and means that they effectively decay more rapidly when the rate at which ants enter the system is high.

The Reward Function

The task that we set our system is substantially simpler than Pask's. We want the ants to arrive at the bottom within a specific range of columns, 19 to 31 inclusive (twelve columns out of the 50 in the system). When an ant hits the target we increase the rate at which ants enter the system. Since each ant lays down the same amount of pheromone (500 units in total) the rate at which ants arrive is proportional to the amount of pheromone added to the system, and therefore limits the total strength of trails that can be maintained in the system.

The details of the scheme used for the presented results are as follows:

1. Let the score for iteration i be $S_i = 1$ if the ant arrives at the bottom of the grid within the target interval, 0 if it misses.
2. This value is smoothed out in time slightly using a leaky integrator: Let the reward value $R_i = R_{i-1} + (S_i - R_{i-1})/\lambda$. We give the parameter λ the value 2.0 and let $R_0 = 0$.
3. Ants are assumed to arrive at a rate $99R + 1$. This is represented by multiplying each pheromone value by $1 - 1/(495R_i + 5)$ (an approximation to $e^{-5(99R_i-1)}$) to represent a constant decay of pheromone during the variable time period between ants arriving.

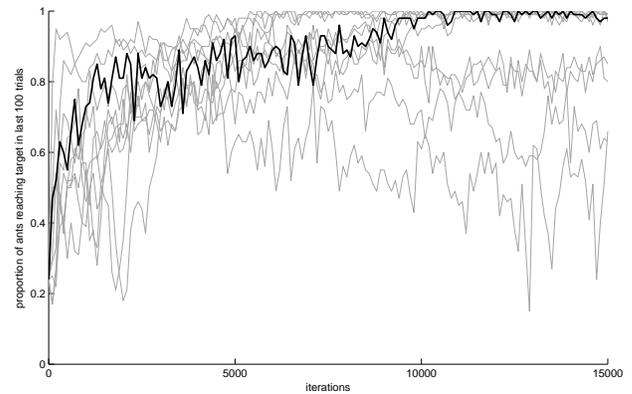


Figure 1: Increase in accuracy over time for ten independent runs of the model. The black line corresponds to the run shown in more detail in figure 2. Each data point represents the proportion of ants which hit the target over a period in which 100 ants are released. The expected probability for an ant to hit the target in the absence of any pheromone trails is 0.24, so the first data point is set to this value for each run.

Experimental Results

Figure 1 shows the proportion of ants hitting the target over time for ten independent runs of the experiment. Four of these systems do not converge on a good solution before the end of the run at iteration 15000 but six of them perform well, converging to a stable state in which very few ants miss the target.

Figure 2 shows the positions of the trails over time for one such run. One can see that the system converges fairly rapidly to a state in which there is a fairly strong trail leading to the target, which has a catchment area that catches almost all the ants entering the system. However, all trails are rewarded for this, not just the ones that lead to the target. This allows several “parasite” trails to persist, which benefit from the ants reaching the target but do not contribute towards it. These are less stable than the one which reaches the target and are eventually out-competed by it.

The four systems which do not converge to a good solution have stronger, more established parasite trails. Since more established trails fluctuate more slowly it can take a very long time for these to decay. Note that although these replicates have not converged on a perfect solution, all but one of them consistently do better than the 25% that would be expected in the absence of any reinforcement effect: they have attained solutions which work but are not perfect. The possibility of getting stuck on ‘local’ optima is perhaps another thing that adaptive processes have in common with evolution by natural selection.

Discussion

It is important to be clear about the relationship of our model to the subjects under discussion, namely Pask's experiment

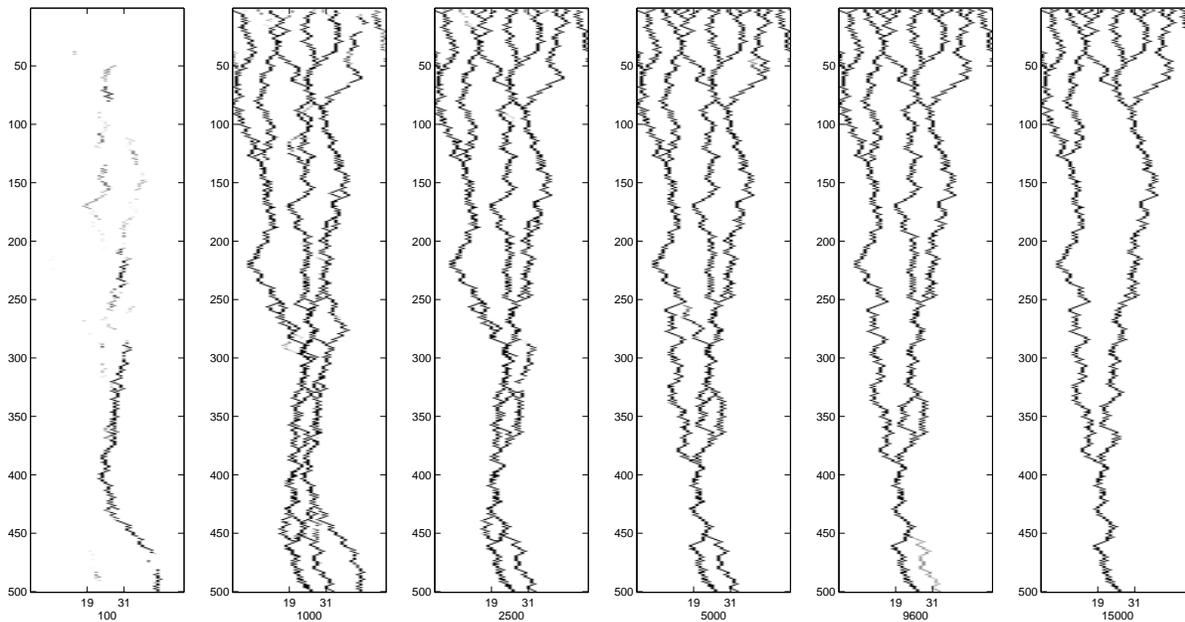


Figure 2: Snapshots of the pheromone trails after 100, 1000, 2500, 5000, 9600 and 15000 ants have passed through the system. The target is the marked interval between columns 19 and 31 at the bottom of the grid. After 100 iterations there is only a weak trail which does not lead to the target. By iteration 1000 a stable trail to the target has been formed which fans out at the top into a large catchment area, but it also supports a number of ‘parasite’ trails which do not hit the target. These gradually disappear, with the last fading out at around iteration 9600. After this almost all the ants hit the target (see figure 1). The system then changes very little until the end of the run at iteration 15000.

and adaptive growth processes. Our claim is that our model and Pask’s device operate according to a common principle, the adaptive growth mechanism that we have described. Our model is intended as a simple instantiation of an adaptive growth process rather than as a direct model of Pask’s ear, although loosely speaking the pheromone trails can be seen as a metaphor for the metallic threads in Pask’s ear, with the rate of input of ants taking the place of electric current.

Comparison to Ant-Colony Methods

Our model is not intended to compete with ant colony optimisation methods as it does not have the same purpose, but since it bears some similarity to them it is worth discussing the how our system differs and the reasons for taking our approach.

It might appear that our system has several disadvantages compared to a more traditional ant-based system in which a successful trail is rewarded. We can see from figure 1 that convergence to a solution is slow and not guaranteed to occur. However, the purpose of our model is explanatory. The threads in Pask’s experiments were not selectively reinforced. Instead he rewarded the whole system with an increase in current flow, and our claim is that our model captures the way in which this directed the system’s growth towards a solution. Moreover, the solution found by Pask’s

electrochemical device does not consist of a single thread; it is a complex solution which requires the co-operation of multiple threads performing a variety of tasks. Rewarding “successful” threads in this context would not make sense, since it is not in general possible to determine which threads are contributing to the solution and which detract from it (see the discussion of the credit assignment problem above).

In our system the task is not to create a single path leading from the top of the grid to the target area, but to create a network of threads which funnels ants from all positions on the top of the grid towards the target area. There is some similarity between this and the system-level nature of the solution found by Pask’s device. Our hope is that with a better understanding of adaptive growth processes it will be possible to design systems, either *in silico* like our model or in physical substrates, which can solve more complex tasks, forming solutions which equal or surpass the sophistication that Pask was able to achieve.

Implications for Adaptive Growth Processes

It seems reasonable to call the growth process in our model adaptive because it shares with Pask’s ear the property that structures which contribute towards performing the task (and thus increasing resource availability) are more stable and out-compete structures which do not help perform the task.

Our computational experiment thus demonstrates that adaptive growth is a general phenomenon, rather than something which only occurs in the specific electrochemical environment of Pask's experiment.

However, adaptive growth does not occur in all possible substrates and the parameters and reward function of our model had to be chosen from the right ranges in order for the phenomenon to occur. For instance, the subtraction of d from the pheromone level in each cell for each ant that enters the system seems to be important for adaptive growth to occur. Without it stable structures that achieve the funnelling task do form, but they spontaneously collapse much more readily, reforming again soon after. Perhaps this is because parasitic structures can grow too easily, diminishing the resource supply and destabilising the whole system. It appears in this case as if the system performs a random walk between more and less stable structures, spending a high proportion of its time in a state containing a stable structure simply because those states change more slowly. But the subtraction of d on each iteration seems to provide a ratchet effect, making transitions from more to less stable structures very unlikely.

Another factor which seems important is that the growth process is fast but that the decay process is slower for more established structures. If the two processes took place on the same time scale then it would be harder for structures to persist on long time scales. The difference in time scales between new and old structures means that it's very unlikely for the whole structure to unravel by chance. If the resource availability drops it is more likely that a new structure will be dissolved than an old one, giving a trial-and-error quality to the growth process.

It also seems important that the reward function is modulated on a similar time scale to the fluctuations in the structure. This is the reasoning behind applying a leaky integrator to our reward function. A good substrate for an adaptive growth process might therefore be one which exhibits fluctuations over a wide range of time scales.

There is clearly a long way to go before we have a full understanding of why Pask's ear works, and of adaptive growth processes in general. Our example is simple and illustrative but we anticipate that more elegant and successful examples will be found which are capable of adapting to more involved tasks, leading to the possibility of practical application.

Conclusion

We have introduced the notion of an adaptive growth process in order to explain the results of Pask's electrochemical experiment. This allows us to see the enormous potential of his result: in terms of practical applications, adaptive growth processes could be used to produce control systems that can adapt to new tasks and even adjust their level of complexity when necessary. The idea may also be important for our un-

derstanding of biology since adaptive growth processes may play a role in development.

We have presented an illustrative computational model in which an adaptive growth process occurs, demonstrating that adaptive growth is a general phenomenon and paving the way for a better understanding of the circumstances under which adaptive growth can occur.

References

- Lotka, A.J.: Contribution to the energetics of evolution. *Proc. Nat. Acad. Sci. USA* **8** (1922) 147–151
- Pask, G.: Physical Analogues to the Growth of a Concept. *Mechanisation of Thought Processes: Proceedings of a Symposium Held at the National Physical Laboratory*. London: H.M.S.O. (1958) 879–922
- Pask, G.: The Natural History of Networks. Yovits, Cameron (eds), *Self-Organising Systems*. New York: Pergamon Press (1960) 232–263
- Pask, G.: A Proposed Evolutionary Model. Van Foerster, Zopf (eds), *Principles of Self-Organisation*. Pergamon Press (1961) 229–254
- Cariani, P.: To Evolve an Ear: Epistemological Implications of Gordon Pask's Electrochemical Devices. *Systems Research* **10(3)** (1993) 19–33
- Bird, J., and Di Paolo, E. A.: Gordon Pask and his maverick machines. P. Husbands, M. Wheeler, O. Holland (eds), *The Mechanization of Mind in History*, Cambridge, MA: MIT Press (2008)