

# THE ORGANISM AS A DYNAMICAL SYSTEM

Peter T. Saunders  
Department of Mathematics, King's College,  
Strand, London WC2R 2LS, England.

(Published in *Thinking about Biology* (SFI Studies in the Sciences of Complexity, Lecture Notes Vol. III, F. Varela & W. Stein, eds). Addison Wesley, Reading, 1993, pp 41-63.)

## 1 Introduction

One of the most characteristic properties of the developmental process is that it is stable. An embryo does not need an absolutely perfect environment and it can survive many small disturbances and even some large ones. Two embryos do not have to be clones to turn into very similar adults. The stability of development is, however, not just the simple sort of stability that we observe in such familiar examples as a ball at the bottom of a cup. An embryo that is perturbed will not return to the state that it was previously in. If it can recover at all, it will continue to develop, eventually reaching more or less the state it would have attained had it been left alone. What is stable is not the state of the embryo at any one time, but its pathway of development.

In biology, the property of returning to the state a system was in before it was disturbed is called homeostasis, from the Greek words meaning similar and standing. C.H. Waddington, one of the first to stress the role of stability in development, introduced the words homeorhesis (similar flow) to describe a system which returns to a trajectory and chreod (necessary path) for the trajectory itself [13, 15, 17]. He also used the term canalization to describe the property that development typically can proceed to one or more of a restricted number of alternative end states rather than to a broad spectrum. And as an aid to understanding the role of these phenomena in development and evolution, he devised the epigenetic landscape, a metaphor which is illustrated in Figure 1.

Waddington arrived at his view from his experience in embryology. Homeorhesis, chreods and canalization describe features which developing organisms do possess, and recognizing their role leads to insights into both development and evolution. It is, however, possible both to see both why existing developmental systems have these properties, and also to realise it would be hard to imagine a developmental system that did not. The properties arise from the fact that organisms, whatever else they are, are dynamical systems. Now an arbitrary system may or may not be stable. If it is not stable, it will not persist, and we probably would not think of it as a system. Consequently a typical property of those entities that we recognize as systems is that they are stable.

When applied to systems in general, this argument is not very enlightening. It tells us more about how we view things than about how things are. But an organism, and in

particular its developmental system, is not just any kind of system. It is a complex, non-linear dynamical system. And such a system, if it is stable at all, generally has not just stability about a single equilibrium point but the richer kind of stability that Waddington identified in development. This allows us to replace the weak statement “The systems we observe are stable” by the much stronger “The complex non-linear systems we observe, and this includes the developmental process in organisms, have a special kind of stability which gives them a number of important properties including those depicted in the epigenetic landscape.” It may be surprising that there are organisms, but once we know that there are, we should not be surprised to find that they have these properties.

While the qualitative study of dynamical systems goes back at least to Poincaré at the turn of the century, it is only comparatively recently that it has become a major focus of mathematical research, largely on account of advances in topology and in computing. The aim of this chapter is to show how modern techniques, even just the modern point of view, can lead to progress in directions that Waddington first indicated many years ago.

## 2 The Epigenetic Landscape

Waddington imagined the developmental system as a mountainous terrain (Figure 1) stretching from the heights, where the process begins, to the lowlands below, where it culminates. The valleys represent possible developmental pathways. The precise shape of the landscape depends on a network of guy ropes beneath it; the ropes stand for the effects of the genes and the complexity of the network is to remind us that while genes certainly affect development they do so in a very complicated way.

A ball rolls down the landscape, and the path that it follows stands for the actual development of the organism. An environmental perturbation is represented either by the ball being pushed up the side of the valley or by a small temporary change in the shape of the valley itself, as if the landscape were a tent and someone had kicked it. In either case, if the disturbance is not too great, then the ball returns to the valley bottom not at the point where it was disturbed but somewhere further down. The system thus exhibits homeorhesis.

If the ball is disturbed so much that it is forced out of the valley, it will probably reach a dead end, but it is also possible that it will pass over a watershed and then continue down a different valley just as stable as the original one. This too illustrates a property which is actually observed: sometimes when an embryo is perturbed it neither dies nor returns to normal development but switches to an alternative pathway which leads to the production of a viable, though significantly different organism.

A mutation is represented by a change in the tension or position of a guy rope which may (though on account of the complexity need not) bring about a permanent change in the shape of a valley. Most of these will be relatively minor and the ball will generally return to its original course.

Mutations are most likely to have significant effects if they disturb the landscape near where one valley divides into two. At such a point even a small alteration in the topography can be enough to send the ball down a different path. Thus a small genetic change can bring about a large change in the phenotype, and without necessarily affecting the action of any other genes.

The model also suggests that environmental perturbations and mutations can have simi-

lar effects: they are simply different ways of diverting development into the same alternative pathway. This is observed in the common phenomenon of phenocopying. If, for example, genetically normal *Drosophila* embryos are immersed for a short time in ether vapour, some of them will develop into adult flies that resemble the *bithorax* mutant [3].

Phenocopying is largely ignored by evolutionists, who seem content to treat everything between the genome and the adult as a black box whose properties can safely be disregarded. It is, however, very important once we start thinking in terms of dynamics. Its significance for the epigenetic landscape is that it provides a clear example of the existence of alternative pathways, i.e. that the neighbouring valleys in the landscape really do exist. For if they did not, it would be very hard indeed to explain how it is that so often a mutation and an environmental perturbation can have the same effects.

### 3 The Epigenetic Landscape and Evolution

The stability of development is very important for evolution. In the picture, the valleys are deep enough that neither minor changes in topography nor small random disturbances to the ball are likely to divert it from its normal course. Since the complexity of the network of guy ropes means that changing the position or tension of any one rope is unlikely to alter the landscape very much, most such changes will have little or no influence on the end result. The majority of exceptions will be near the end of the process, where the sides of valleys are flattening out and where the ball may not have enough time to return to its original trajectory.

Many perturbations, whether mutational or environmental, have little or no effect on the phenotype. There is typically very great genetic variability in populations which have very nearly identical phenotypes. The model (or, if you prefer, the recognition of the role of chreods) makes it easy to understand why this should be so.

Like most theories, the model has no difficulty with small evolutionary changes, but we have to ask what it suggests about large ones. In principle, of course, they might occur by long sequences of small changes. This would require a succession of small deformations of the landscape, each moving the end point a little bit until it had reached a significantly different position. While this is possible in principle, it seems unlikely. If pulling a little on a rope shifts the landscape in the appropriate direction at first, after a while the tensions from the other ropes to which it is linked will stop it from continuing the same effect. And at the same time, the complexity of the network means that other parts of the landscape may also be affected, and in ways that were not intended.

Again, what the model illustrates are real properties. Even artificial selection seems inevitably to run up against limits. It is also generally accepted that both pleiotropy (one gene affecting more than one character) and genetic linkage (the fact that genes that are physically close together on a chromosome tend to be passed on together) will oppose strong directional selection. A sequence of advantageous changes in one character is likely also to bring about a sequence of random changes in others, and if this is carried on for some time the net effect will almost certainly be deleterious.

But we have already noted that a small change in the landscape can have a significant effect on the organism if it occurs just where it can divert the ball down a different path. No other part of the landscape would have to be affected significantly, because any concomitant

changes would also be small and all that is needed is that none of them is at a critical point. What is more, since it is only a matter of a diversion into an existing path, not the creation of a new one, we would expect that any one of a number of different small alterations in the network of ropes would do.

The model thus suggests not only that a single mutation can bring about a large phenotypic change but also that different single mutations can have the same effect. This could be because different mutations led to the production of the same protein, though we shall see later that this is not the only way. This helps us to understand how evolutionary change can occur without requiring that the same mutation occur more or less simultaneously in a large number of organisms. (As Haldane [2] showed many years ago, a mutation that occurs in only one individual in a large population is unlikely to survive, even if it confers a significant selective advantage.)

The model therefore predicts that large changes will happen rapidly and will probably not be related to any minor ones that are going on at the same time. Thus the mode of evolution suggested by the epigenetic landscape is precisely that of punctuated equilibria.

This is a very significant difference between the genetic and epigenetic views of evolution. From Darwin to the present, most evolutionists have insisted that evolution is gradual. Some are now prepared, albeit reluctantly, to accept that large changes might occur, and they see them as caused by mutations in regulatory genes that affect the action of several other genes. Such mutations can no doubt occur, but it is very hard to see how the effect of all the changes they cause will be to produce a viable organism, let alone one that will be in any sense fitter than the normal one. The genetic theory does not predict large changes. At best it suggests how they might possibly be accommodated within the theory.

## 4 The Dynamics of Development

The epigenetic landscape illustrates remarkably well some important properties of developmental systems. It does not, however, explain why they should all have these properties, when there are so many differences between organisms or even between different regions of a single complex organism. The answer is that the properties are common not just to developing organisms but to most non-linear dynamical systems, certainly to those of any complexity.

Waddington was aware that one should think of development in dynamical terms and he even wrote down some equations as an indication of how this might be done. In fact, he was closer to the modern approach when he drew the epigenetic landscape. For we can define a dynamical system as a manifold with a vector field defined on it. In other words, we describe a dynamical system by specifying the complete set of possible states, the phase space, and then providing rules that tell us at each point in the space where to go next. The rules can be given as differential equations (as they will be in the rest of this chapter) but they could also be, for example, difference equations or cellular automata.

In the case of the epigenetic system, the manifold is the Euclidean plane. The vector field is provided by the landscape: the slope of the landscape at any point indicates the direction in which the ball will be accelerated. If we think of the system in more than two dimensions, as we should, the manifold is  $\mathbb{R}^n$ ,  $n$ -dimensional Euclidean space, but the situation is otherwise much the same, except that it is harder to visualize and impossible to

draw.

Actually, it is not quite correct to say that the epigenetic landscape is a dynamical system. Figure 1 is not meant to depict a particular developmental process, only to illustrate some typical features of development. So it would be better to say that the epigenetic landscape represents a class of dynamical systems which share a number of important properties which are typically found in developmental systems. The mathematical problem is to determine the properties of this class.

In the rest of this chapter it will be assumed that many processes of development can be modelled by one particular kind of dynamical system, differential equations. This is not a difficult hypothesis to defend, and besides, it is part of practically all conventional modelling in development. The results will mostly be consequences of the assumption that the differential equations are non-linear and therefore have many properties that are not found in the linear systems with which we are more familiar.

## 5 Nonlinear Differential Equations

It should not come as a surprise that non-linear differential equations are significantly different from linear ones. Remember what happens with algebraic equations. If  $a$  and  $b$  are real and  $a$  is not zero, then the equation

$$ax + b = 0$$

has the unique solution  $x = -b/a$ , which is a real number. Even the more complicated case, a system of  $n$  linear algebraic equations in  $n$  unknowns, still usually has a unique solution, and any solutions must be real.

On the other hand, the quadratic equation

$$ax^2 + bx + c = 0$$

with  $a, b$  and  $c$  all real and  $a$  not zero, has two roots, and these can be complex numbers. So the transition from linear to non-linear introduces multiple solutions and a completely new kind of solution as well.

A quadratic term is about the smallest amount of non-linearity you can have, so you might ask what happens if we make the equation much more non-linear, with higher order terms or even functions that are not polynomials. The answer is: not a lot. There may be more roots, and they may be harder to find, but that's about all. Even if the coefficients in the equations are complex, nothing worse happens: there is no new kind of number lurking beyond complex. The most important differences between linear and non-linear equations appear right away. Once we understand the quadratic case, we know the most important consequences of non-linearity. (To the sort of accuracy to which Nature works, we may ignore the distinction between rational numbers on the one hand and irrational and transcendental numbers on the other.)

Things are much the same with ordinary differential equations. A linear ordinary differential equation, or a set of them, typically has a unique critical point. If this is stable, the system remains at equilibrium or, if disturbed, returns to it. If it is unstable, then the system does not persist. And that is pretty much the full repertoire of a linear system.

The behaviour of non-linear systems is much more interesting. They can have more than one equilibrium point, so that depending on the initial conditions or perturbations during their course they can end up in different final states. These states typically form a discrete set, not a continuum. Besides stable points, however, they can also have stable trajectories. If the system is moving along one of these trajectories and is perturbed not too far away from it, it will return. Because there can be more than one stable trajectory, if a system undergoes a large disturbance it may not return to its original trajectory but may be drawn to another stable one instead.

If a system has more than one stable equilibrium or stable trajectory, and almost all but the very simplest stable non-linear systems are bound to, we may expect transitions from one to another. These will typically happen on a time scale that is short compared with the rest of the behaviour of the system. Relatively small changes in parameters can, by causing attractors to disappear, bring about large changes in the state of the system. Thus large effects do not have to have large causes, as they generally do in linear systems. What is more, in non-linear systems different causes can have the same effect. In particular, when a large change is possible, it can usually be initiated by altering the value of any one of a number of different parameters. This is very important for evolution, because it means that not only are large changes possible, they do not depend on one particular mutation nor even on a small set of mutations each with very much the same effect.

Another characteristic of non-linearity is that we cannot add solutions: the whole may be greater than the sum of the parts or less, but it is seldom the same. This too is familiar from elementary algebra: for example, the square of the sum of two numbers is not just the sum of the squares of the same numbers.

All these properties of non-linear systems are well known, and have been for many years. What is more, as with algebraic equations, only a small amount of non-linearity is needed to make them appear: for instance, the famous Lorenz attractor has only two non-linear terms and even they are only products of two variables [12]. Unfortunately, most of our intuition is based on linear systems because they are easier to analyze and are therefore more common in undergraduate courses. Even when modelling leads to non-linear equations, the first thing we do is generally to linearize the system so that we can study its behaviour near the critical points.

Fortunately, the ready availability of powerful computers is bringing about a change in the situation. Mathematicians are devoting a great deal of attention to non-linear systems and even non-mathematicians are becoming acquainted with their remarkable behaviour. When you look at one of the beautiful fractals that can be generated from simple mathematical formulae [7], you are seeing just one of the many consequences of non-linearity.

There is nothing mysterious about the properties that Waddington identified as typical of developmental systems: they are likely to be found in any complex system which can be modelled by non-linear differential equations. There are also more such properties than Waddington found. For while non-linear differential equations – and so the systems that can be modelled using them – typically have all the properties of development that are illustrated in the epigenetic landscape, the converse is not true. Ingenious though the picture is, it is only a picture and cannot be expected to capture all the properties of non-linear systems.

## 6 Multiple Speciation

One of the shortcomings of the epigenetic landscape is that it is drawn as a two-dimensional surface. It has to be, because that is the most that can be shown in a picture. Waddington acknowledged that this was a simplification, but it actually makes much more difference than he thought, because two dimensions is a very special case in dynamics. This is largely because we are concerned with trajectories, which are one-dimensional, and there are many things that are true only when the dimension of what you are studying is precisely one less than the space that it is in. One important difference is that in the landscape each valley has only two neighbours, but in higher dimensions a path can have many neighbours: imagine a multi-core electric cable. This means, for instance, that there can be a very large number of dead ends and still a few viable pathways, which is not obvious from the picture.

Note how the crucial difference again appears at once. It does not matter very much how many dimensions there are, as long as it's more than two. Another peculiarity of differential equations in two dimensions, incidentally, is that the only attractors are points and limit cycles. In three or more dimensions there can also be strange attractors, i.e. chaos.

If developmental pathways can have many neighbours, we may ask if multiple speciation can occur. The dynamical systems approach suggests that it should. To see why, it is perhaps easiest to think of a simple physical analogue. If you gradually increase the loading on a pillar, for a long time nothing much happens. This is convenient, because otherwise your house would sag appreciably as soon as you moved the furniture in.

If you increase the load beyond a certain critical value, however, the pillar will buckle. The direction in which this happens depends on either an asymmetry in the pillar which makes it weaker in one direction than in others, or else an imbalance in the load. Sometimes the bias is obvious, as when a lumberjack cuts a notch to ensure that a tree falls in the right direction, but often it is not, and then it can be impossible to predict in which direction the failure will occur.

Now imagine a number of symmetric pillars, each carefully made to the same specification and with no obvious flaws, and suppose each is given the same, gradually increasing load. For a while, the pillars will appear not to respond. Eventually, however, and more or less together, they will all buckle. But they will not all buckle in the same direction, because the minute biases that determine the direction of buckling will be different. When the stability of a system breaks down, previously unimportant distinctions can become crucial.

In the same way, the stability of the epigenetic system, which is necessary for normal development, means that individual organisms will develop in the same way even though there are considerable genetic differences among them. What is more, so long as evolution involves only minor phenotypic changes, canalization ensures that most individuals change in the same way. For speciation to occur, the stability has to break down, and when that happens these previously unexpressed genetic differences may come into play, diverting development into different developmental pathways. We may therefore observe the almost simultaneous appearance of a number of new forms. This will not always happen, but we should not be surprised when it does.

The idea that genes that previously had no effect on the phenotype can later come into play is not new, but it is usually seen as a matter of silent genes being turned on. Here we are supposing that the genes are already on, but that the stability of the epigenetic system has been preventing differences in them from affecting the phenotype. When this

stability breaks down, the previously unimportant (but not unexpressed) genes now have the opportunity to influence development, since the relatively unstable system is susceptible even to small genetic differences. In this way a number of different genes can become significant simultaneously without having either to appear or first be expressed simultaneously.

## 7 Direction in Development and Evolution

We now demonstrate that where transitions between distinct states can occur, they are more likely to happen in one direction than in the other. This lack of symmetry, which cannot be inferred from the picture, arises from the details of a physical or chemical process, but can be passed up through the dynamical system to affect processes with quite different time scales, including ultimately evolution.

To make the argument easier to follow, and to show that we do not have to postulate a complicated and implausible mechanism, it is convenient to work with a particular equation. The result does not, however, depend on this choice, but is true for almost any model which can produce a sharp transition between states.

The model that will be used as an illustration was proposed by Lewis, Slack & Wolpert [4], who wanted to show how a sharp frontier could form in a region of tissue. The states of cells in the region were supposed to be specified by the concentration of a gene product,  $g$ , which is activated by a “signal substance”  $S$ , and the rate of change of  $g$  was given by the equation

$$\frac{dg}{dt} = K_1 S + \frac{K_2 g^2}{K_3 + g^2} - K_4 g, \quad (1)$$

where the  $K_i$  are all constants. Lewis, Slack and Wolpert set  $K_2$  and  $K_3$  equal to unity and  $K_4$  equal to 0.4. Without loss of generality we can choose the units of  $S$  such that  $K_1$  is also equal to unity.

Suppose that both  $g$  and  $S$  are initially zero and that  $S$  is then gradually increased. Then  $dg/dt$  will be positive, so  $g$  will increase as well. If the increase in  $S$  is slow enough,  $g$  will always be close to the equilibrium value (for the current value of  $S$ ) which can be found by solving the equation  $dg/dt = 0$  for  $g$ . With the given values for the constants this equation is

$$2g^3 - 5g^2(1 + S) + 2g - 5S = 0, \quad (2)$$

and for  $S = 0$  it has three real roots: 0, 0.5 and 2. There are thus three possible steady states, and it is not hard to show that those at  $g = 0$  and  $g = 2$  are stable and that at  $g = 0.5$  is unstable.

Thus  $g$  will remain zero until  $S$  begins to increase. When  $S$  is greater than zero but small, Eq.(2) still has three real roots and so there are still three equilibrium points, two stable with an unstable one in between. The smallest equilibrium value of  $g$  increases with  $S$ , which means that as  $S$  increases slowly, so does  $g$ .

If, however,  $S$  is increased above a critical value  $S_c$ , which is approximately 0.0418, a significant change occurs. Eq.(2) now has only one real root, together with a complex conjugate pair. The two steady states (one stable, one unstable) corresponding to smaller

values of  $g$  have coalesced and disappeared. The system will therefore move rapidly to the higher equilibrium, i.e. there will be a rapid increase in  $g$ . Thus at this point a very small change in  $S$ , from just below  $S_c$  to just above it, will cause a large change in  $g$ .

This gave Lewis, Slack and Wolpert what they required, because if there is a smooth gradient in  $S$  throughout a region, at the position where  $S$  takes on the value  $S_c$  there will be a discontinuity in  $g$ , i.e. a definite frontier between two sub-regions. Thus the continuous gradient has given rise to a sharp division. Here we are not so much interested in the formation of a frontier as in an abrupt change in evolution, but the model will provide this too, if we imagine that  $S$  has the same final concentration throughout the region but that a mutation increases this value from below  $S_c$  to above it, or vice versa. The principle is the same, but the important variable is time rather than space.

It turns out, however, that there is more to the model than first appears [10]. Eq.(1) has more properties than we have seen so far. What is more, these properties are likely to be found in almost any system which is capable of producing a transition from one steady state to another.

We know this as a result of catastrophe theory, which, in its non-controversial role as a theory about classes of systems of differential equations, tells us that a system that can produce two steady states should have at least two parameters. A model with only one parameter is structurally unstable, which means that almost any other model only very slightly different from it makes different predictions. If we believe that Eq.(1) with all but one parameter fixed is precisely the right model for a particular process, that is one thing, but it cannot tell us how a typical process of this kind will behave.

Fortunately, it is both necessary and sufficient to allow two parameters to vary. Here we shall choose  $K_3$  (which we shall write simply as  $K$ ) because varying a saturation constant seems physically plausible. Allowing the other parameters to vary as well would not change the general pattern of the behaviour, only the details. Other, quite different models, providing that they were not unnecessarily complicated, would also exhibit the same general pattern. We shall see later on that the chief result of this section applies to more complicated models as well, even though Eq.(1) does not describe their entire repertoire of behaviour.

With this change, the equilibrium condition is

$$2g^3 - 5g^2(1 + S) + 2Kg - 5KS = 0. \quad (3)$$

Figure 2 is a diagram of the ‘‘control space’’ for the system, i.e. the  $S - K$  plane. The idea is that the system is controlled by variations in the parameters  $S$  and  $K$ , and responds by adjusting the gene product  $g$  to an equilibrium value. This value will naturally depend on  $S$  and  $K$ . Since Eq. (3) is a cubic, it has three roots, but they are all real only for certain combinations of the parameters. For others, there is one real root and a complex conjugate pair. Hence for some values of  $S$  and  $K$  there are two stable equilibrium values of  $g$  (and one unstable one in between) whereas for others there is only one. The cusped curve in Figure 2 is the bifurcation set, i.e. the boundary between the two regions. (Those familiar with catastrophe theory will have recognized Eq. (3) as a cusp catastrophe, though not in its canonical form [8].)

In general, a small change in either  $S$  or  $K$  will cause only a small change in  $g$ . The exceptions are when the change takes the system from the region of the control space in which there are two equilibria to that in which there is only one. If the equilibrium that the




FIGURE 2 The control space for Eq. (3), showing the bifurcation set. The regions in which only the A phenotype, either the A or B, or only the B phenotype are possible are marked by A, AB, and B, respectively. The lines  $K = 1$  and  $S = 0.0418$  are relevant to the two models described in the text.




FIGURE 3 A qualitative projection of the control space for Eq. (4), showing the bifurcation set and control trajectories for  $L = 0.55, 0.60, 0.65, 0.75$  as  $S$  rises from 0 to 0.5 or above. The letters A, B, and C indicate the possible phenotypes as in Figure 2. All three are possible in the unlabelled diamond-shaped region in the center of the figure.

system is in is the one that disappears (by coalescing with the unstable equilibrium to form a saddle point) then there will be a sudden change in  $g$  as the system moves rapidly to the other equilibrium.

Suppose that  $S$  is initially zero. Then  $g = 0$  is a (stable) equilibrium, so if the gene product is initially at zero concentration, it will remain there. If  $S$  is then increased slowly enough that the system is not disturbed appreciably from steady state,  $g$  will also increase slowly. This will continue until the increase in  $S$  takes the system across the right hand branch of the bifurcation set: for  $K = 1$  this is at  $S = 0.0418$ . When that happens, the low value equilibrium disappears and  $g$  will increase rapidly. (For the values used in this illustration, the change is from  $g = 0.227$  to  $g = 2.165$ .)

If  $S$  is then decreased, nothing significant will happen as the trajectory crosses the right hand branch because the high valued equilibrium is still present. The sudden jump back to a low value occurs only when the trajectory crosses the left branch of the cusp, which for  $K < 1.563$  would require  $S$  to be negative. Hence for all values of  $K$  below 1.563, even if  $S$  returns to zero the concentration of the gene product  $g$  will remain at a significantly non-zero value; for  $K = 1$ , this is 2.0.

This is an example of the very important phenomenon of hysteresis. When a system can move from one state to another, even if the reverse transition is possible the system will not simply retrace its path in the opposite direction. As the name hysteresis (from the Greek word for delay) suggests, systems tend to remain close to whichever end state they happen to start in. Because of this, even reversible transitions are typically not symmetric, and, as we shall see, can contribute to irreversibility.

We now consider two different ways in which the model we have been describing might be involved in development. For simplicity we will suppose that whether  $g$  has a low value or a high one is directly observable in the phenotype and we will refer to individuals with low and high values of  $g$  as A and B, respectively.

## 7.1 MODEL I

Suppose that during development  $K$  is fixed and  $S$  rises slowly from 0 to a maximum value  $S_m$ . Both  $K$  and  $S_m$  are assumed to be under genetic control. Then for  $K$  less than about 2, combinations of  $K$  and  $S_m$  that lie in the regions marked A or AB will produce A individuals (since  $g$  will never be forced away from the low valued equilibrium), but those that lie in the region marked B will produce B individuals.

Now suppose that mutations alter the values of  $K$  or  $S_m$ . Any mutation which did not take the values across the right hand branch of the cusp, i.e. the boundary separating region AB from region B would have no observable effect on the phenotype. On the other hand, a mutation which caused only a small change in either parameter but which did take the system across that boundary would produce an obvious effect.

Thus, for example, imagine that  $S_m$  is fixed at 0.0418 and  $K$  varies from 0.5 to 2.0. For all values of  $K$  less than 1.0, the result would be B, and for all values of  $K$  greater than 1.0 the result would be A. Whether the change in  $K$  occurred in a single step or as an accumulation of many mutations, the large change in  $g$  and therefore in the observable phenotype would occur abruptly. So the mathematical model, like Waddington's picture, predicts punctuated equilibria.

Suppose that after  $K$  and  $S$  have reached their normal values there is a change in the concentrations of one or both of them. If  $K < 1$  then there will be no net effect. If  $K > 1$ , however, so that the normal phenotype is A, then a temporary reduction in  $K$  or increase in  $S$  which carried the point across the right hand branch of the cusp would cause  $g$  to move to the high equilibrium. Even if  $K$  and  $S$  returned to their original values,  $g$  would remain at the high equilibrium, and so the phenotype would be an B. This phenocopying is most likely to happen for  $K$  not very much greater than 1.0, because as can be seen from the diagram, as  $K$  increases so does the size of the perturbation required to have a significant effect.

If the gene product  $g$  is perturbed, it will return to an equilibrium value for the given values of  $K$  and  $S$ . If  $K < 1$  there is only one equilibrium and so the organism will still be a B. If  $K_4$  lies between 1.0 and 1.85 there are two equilibria. The value of  $g$  will normally be at the lower equilibrium, but if  $g$  is increased too far it will not return to this equilibrium. Instead, it will stabilize at the higher one, i.e. the organism will become a B instead of an A. As  $K$  increases the two equilibria move further apart, so that the transition becomes less likely, and finally when  $K > 1.85$  the higher equilibrium disappears and the phenotype must be A.

Thus as  $K$  increases, there is a range within which the phenotype must be B, a range within which the normal phenotype is A but B phenocopies can occur, and finally a range in which the phenotype must be A. Phenocopying never occurs in the direction B to A.

## 7.2 MODEL II

Model I follows Lewis, Slack & Wolpert's original idea, that  $S$  rises from zero to a final value  $S_m$  where it remains. Saunders & Ho [10], however, suggested that a more reliable way to create a boundary would be to establish a gradient in  $K$  and then have  $S$  rise to a large but not precisely specified value, say 0.1 or thereabouts, and then fall back to zero. The boundary would then occur at a particular value of  $K$ . The advantage would be that since for any  $K$ ,  $g$  remains zero as long as  $S = 0$ , there would be plenty of time to set up an accurate gradient in  $K$ . To set up an accurate gradient in  $S$  is harder, because as soon as  $S$  starts to increase, so does  $g$ . Here we are not concerned with boundaries, but it is still of interest to see how this different dynamic behaves.

Because the maximum value attained by  $S$  is no longer critical, we may assume that the only variable under close genetic control is  $K$ . If there are no perturbations, then all individuals with  $K > 1.563$  are A and all those with  $K < 1.563$  are B.

Now suppose that  $K$  undergoes a perturbation, which we again take to be a small and temporary change from its normal value  $K_o$ . If this happens before the rise and fall in  $S$  it will have no effect because so long as  $S = 0$ ,  $g = 0$  is a stable equilibrium for all values of  $K$ . If it happens during or after the rise, however, the equilibrium reached when  $S$  falls back to zero will be that corresponding to  $S_t$ , the temporary value of  $K$ . The only interesting cases are of course when one of  $K_o$  and  $K_t$  is less than 1.563 and the other is greater. If  $K_o > 1.563$  and  $K_t < 1.563$ , as  $S$  falls,  $g$  will remain at the high equilibrium, but when  $K$  returns to  $K_o$  this is no longer available,  $g$  will fall to the low equilibrium, and the organism will be an A, as it should. In the reverse situation however, as  $S$  falls  $g$  will go to the lower equilibrium, and since this does not disappear when  $K$  returns to  $K_o$ ,  $g$  will remain there. So the genetically B organism can be perturbed to an A phenotype.

Alternatively, since for  $K < 1.563$  there are two equilibria but for  $K > 1.563$  there is only one, a perturbation of  $g$  can cause a permanent change in the former case but not in the latter. So again, phenocopying can occur only in the direction B to A. The two models differ in the direction in which phenocopying is possible. This is not important, because if Eq.(1) described a real process in development, either model I or model II would have to apply, not both. What is significant is that there is always an asymmetry.

Like the epigenetic landscape of which it is a mathematical version, the model predicts that phenocopies can occur. It also predicts that during evolution from one form to another, there may be a limited period during which phenocopying is possible, but that the system may later stabilize so that it is not.

The most interesting prediction, and one which cannot be seen in the epigenetic landscape, is that we expect phenocopying in one direction only. We know that we can readily produce a *bithorax* phenocopy from a genetically normal *Drosophila* embryo; this should not lead us to expect that we can produce a phenotypically normal *Drosophila* from a *bithorax* mutant embryo.

When an evolutionary change occurs there may be a period just before the change in which the form that is about to take over appears from time to time as a phenocopy. Alternatively, there may be a period shortly after the change in which the old form sometimes appears as a phenocopy. But we should not expect to see both patterns of behaviour.

Finally, the model illustrates the point made earlier, that the mutation required to cause a large change is unlikely to be unique. Model I has only two parameters,  $K$  and  $S_m$ , and yet any mutation that altered either of them in the right direction would do. And in the original model given by Eq. (1), the same effect could be brought about by a mutation affecting any of four different parameters. The magnitudes of the changes caused by different mutations would not have to be the same because the magnitude of the change in gene product is a property of the dynamic, not the mutation. In the numerical example given above, almost any mutation that changed  $g$  significantly would increase it from about 0.2 to about 2.1.

Dynamical systems typically have more variables than are necessary to produce all the behaviour of the system. In the model we have been discussing here, there are four physical variables but only two that are in this sense mathematically independent. Since it was chosen as an illustration rather than proposed as a mechanism for a particular process, we may expect that in many real situations there would be even more physical parameters, even though the behaviour was in general terms the same.

Mathematicians naturally prefer to work with model systems that have no such redundancy. The first step in analyzing a system is often to identify and eliminate inessential parameters, as we have done here with Eq. (1). The value of this strategy is obvious, but it does tend to give the misleading impression that each mode of behaviour of a dynamical system can be attributed to a single physical parameter. Not every parameter is necessarily capable of causing every kind of change, but the real situation is usually far more complicated than the idealized models in textbooks.

Biologists now recognize that neither “one gene/one character” nor “one gene/one polypeptide” is a valid assumption. Since the polypeptides ultimately act through the parameters of chemical reactions, we can add the third side to the triangle: “one polypeptide/one character” is also false.

## 8 Genetic Assimilation

The beginning of a major transition often poses a problem for Darwinian evolutionary theory. However great the selective advantage of the final result, it can be very hard indeed to imagine what possible use the first step towards it could have been.

A famous example, due to Mivart [6] in 1871, concerns flat fish like plaice and flounder, which swim on their sides near the bottom of the sea and have both eyes on the same side of the head. An eye that pointed downwards would presumably not be of much use, so let us agree that there is a significant selective advantage in having both eyes on top. Let us also suppose that the change from the usual arrangement occurred gradually, by a number of separate mutations. Then each of these mutations must have individually conferred a significant selective advantage, and this is not so easy to see. Above all, what benefit could there have been in having the bottom-facing eye just marginally further forward than before? Yet if there was no benefit, how did the process of moving the eye ever get started?

Darwin's own suggestion [1], incidentally, was that the first stages of the transition "may be attributed to the habit, no doubt beneficial to the individual and to the species, of endeavouring to look upwards with both eyes, while resting on one side at the bottom." This may have been an adequate response at the time, since what is now called Lamarckism was then widely accepted, but most modern evolutionists would not be happy to rely on it.

Suppose, however, that the asymmetric form could occur as a phenocopy, and that each of the mutations not only moves the eye slightly but also increases the probability that the phenocopy will occur. Then each mutation would increase fitness by increasing the chance that the individual that possessed it has both eyes on top. A small fraction of an advantageous trait may be of no use; a small probability of having the trait is another matter.

Whether genetic assimilation plays a significant role in evolution is not known, though Waddington [15] and others have demonstrated it in the laboratory. But the simple example here illustrates in concrete terms how it could come about. Suppose that  $K$  is initially 1.4 so that the normal phenotype is A, and suppose that the B form is fitter. Suppose also that any point mutation that affects  $K$  reduces it by 0.1. A single such mutation would not change the phenotype, and so would appear to have no selective advantage. But it would increase the chance of the individual that possessed it becoming a B phenocopy, and this would give a selective advantage. As a result, we would expect genes that reduced  $K$  to increase and that the B form would eventually replace the A.

Our results suggest that phenocopying will frequently be observed when a significant evolutionary change is about to occur. They also show how genetic assimilation can contribute to irreversibility in evolution, because if it assists a change in one direction it will not help in the other. This could help to explain Dollos's law, which states that a feature that is once lost is unlikely ever to reappear.

## 9 More Complexity

The results of the preceding sections depend on two slightly different assumptions about simplicity. The first was that the transition from one state to another occurred in a way that is relatively easy for nature to bring about. An electric light switch may operate without

a hysteresis loop, but it is an artificial device. What is easy for us is not always easy for nature.

We also assumed that the mechanism was the simplest natural one that can produce a transition between states, and here simplest was used in the sense of producing such a transition and nothing more. The next simplest class are those with three stable equilibria, and as a concrete example we can use the following extension of Eq.(1):

$$\frac{dg}{dt} = S + \frac{ag^2}{K + g^2} + \frac{bg^2}{L + g^2} - Dg \quad (4)$$

This seems a plausible example, but its real justification is that we know from catastrophe theory that Eq.(4), like the simpler Eq.(1), captures the essential behaviour of a very large number of systems, many of them far more complicated and with many more variables.

For most values of the parameters there are only one or two stable equilibria, and so most of the behaviour of this equation is very much like that of Eq.(1). The results of the previous section therefore apply. If, however, we take  $K = 36.6$ ,  $A = 1$ ,  $B = 0.14$ ,  $D = 0.11$  and  $L$  greater than zero but less than about 2.5, then for small values of  $S$  there are three equilibria, which allows more possibilities than before.

Some of these are illustrated in Figure 3. The axes have been neither scaled nor even labelled on the figure because it is to be interpreted only qualitatively. This is necessary because the analogue of Figure 1 would require four dimensions, and so the most that can be shown is a projection. (A full description of the figure and how it is to be interpreted is given by Saunders and Kubal [11].)

There are now three possible equilibria, one at low  $g$ , one at an intermediate value, and one at high  $g$ . These have been labelled A,B,C, respectively. The four curves illustrate what happens for different values of  $L$  as the signal substance  $S$  is increased from 0 to a large value. Because we are concerned only with the sequence of equilibrium states that the system will pass through under various conditions, a figure that is only topologically correct is good enough. If we took the model seriously and wanted to compute the critical values of the four parameters we could always do this starting from Eq.(4).

As in Eq.(1), if  $S = 0$ , then  $g = 0$  is a stable equilibrium, so all four systems start in state A, the low- $g$  equilibrium. What happens after that depends, as before, on the role of  $S$ . If we suppose it rises to a precise maximum value, and remains there, the behaviour is almost exactly as in model I of the previous case. The system will remain in state A until the corresponding equilibrium disappears, which occurs at the value of  $S$  at which the trajectory crosses the part of the bifurcation set that separates the regions marked CA and C. If the normal value of  $S$  is below that value but  $S$  is temporarily raised above it, the system will move to C and stay there; if the normal value is above the critical value and falls below it for a short time, nothing much happens. Unless  $S$  is very small, if  $g$  is at the low equilibrium it can be perturbed to the higher equilibrium, but if  $S$  is large enough so that  $g$  has moved to the higher equilibrium there is no longer a lower one for it to be perturbed to. Hence phenocopying can occur in the direction A to C only.

State B cannot be reached by mutation or by perturbations in  $L$  or  $S_m$ , but it can be reached by a perturbation of  $g$  when  $L$  and  $S_m$  have appropriate values. At such values, the normal state is A, so phenocopying will be in the direction A to B, not C to B. Thus the only phenocopies are A to B and A to C.

The absence of normal B phenotypes arises out of a distinction between mathematical systems and the real phenomena they are being used to model. Large regions of the phase space of a mathematical model are often inaccessible in reality. The most obvious reason for this is that many physical variables cannot take on negative values.

An important consequence of this is that some regions of the phase space which correspond to permitted values of the parameters are for the most part inaccessible because they can only be reached from normal starting points by passing through a forbidden region. In effect, the phase space for the real process may not be connected. It might be, however, that a perturbation could force the system into such a region. This would correspond to forms which can occur only through environmental disturbance, never by normal development or mutation.

As a general point, it is worth noting that while the assumption that a phase space is a smooth, simply connected manifold is seldom stated explicitly, many results depend on it. For example, a number of authors have discussed evolution in terms of fitness landscapes with species striving to reach adaptive peaks, i.e. phenotypes with (locally) maximum fitness. In principle there is no reason why one cannot imagine a phase space for phenotypes. On the other hand, its topology (loosely speaking, its structure) would be very different from that of a Euclidean space. At any point we can move only in a very restricted number of directions, and we cannot go from almost any point to almost any other point by a continuous curve. (Translated into biology: the number of possible variants is much fewer than the number of conceivable alterations in the phenotype, and not every transition from any one phenotype to any other can be accomplished by infinitesimal changes.) That there can be large phenotypic changes implies that there are points in the phase space that are in one sense far apart and yet in another sense close together. We should not assume that just because it is possible to define a space, it, and therefore the system or process it is supposed to model, will have all the nice properties of the Euclidean plane.

Returning to the model, let us consider what happens if we assume that development involves  $S$  becoming quite large, say 0.5, and then falling back to zero, like Model II in the first example. It is then easy to track what happens on the figure. For  $L=0.55, 0.60, 0.65, 0.75$  the final states are C,C,B,A respectively. Temporary changes in  $L$  or in  $g$  can cause permanent transitions from C to B or A or from B to A, but not the reverse. Note that these are all in the opposite direction to those permitted when we assumed  $S$  remained at its maximum value. Thus we have the same result as before: the direction in which phenocopying is possible depends on the details of the process, but it is only possible in one direction.

In the model based on Eq.(1), we could see all the interesting behaviour using obviously plausible values of  $K$  and  $S_m$ . In Eq.(4), on the other hand, it takes a fair amount of trial and error to find parameter values that give anything new. This might be an artefact, but it probably reflects a general rule that even if complex systems do have the potential for complex behaviour, most of the time they act like simpler ones. For example, a complex elastic structure may have many different failure modes, but it will still usually behave in much the same way as the pillars mentioned above, showing little response to loading until the critical load is reached, only then buckling. If, however, the parameters have just the right values necessary to make the critical loads for two failure modes equal, the behaviour will be quite different. The imperfection sensitivity will be very great, which means that the maximum safe load may be much lower than expected, and if the structure is overloaded it

will collapse without warning [14].

While the consequences are unlikely to be as disastrous as if we were building bridges, we too should remember that a complex system can have hidden within it a capability for much more complicated behaviour than it usually displays. The potential is likely to be realized only rarely, but then significant steps in evolution do not occur very often.

## 10 Time Scales

In *The Strategy of the Genes*, Waddington [17] pointed out that a living thing is involved in at least three different types of temporal change, all going on simultaneously. These are, in order of time scale, physiology, life history and evolution. Development is part of the medium time scale, and interacts both with the short term physiological processes that bring it about and also with the long term process of evolution. Biologists, faced with the immense complexity of organisms, have naturally tended to look at the three time scales separately, and physiology, developmental biology and evolution exist as separate disciplines. We might call this a sort of epistemological adiabatic approximation, drawing on a term that is sometimes used in physical science to describe the mathematical decomposition of a dynamical system into sub-systems with different time scales.

This device is often very useful in mathematics, and it is certainly hard to imagine how biology could progress if we had to take absolutely everything into account all the time. The subject is hard enough as it is. All the same, we must not forget that the levels do interact, and in both directions. It is not simply a matter of the slowly varying quantities acting as parameters for the fast processes, of evolution setting the framework for development. Here we have seen how a fast process can profoundly affect the nature of a slow one. The origin of the asymmetry was that the gene product  $g$  started at zero rather than a high value. This created a lack of equivalence between two equilibrium values, i.e. between two developmental states. This in turn meant that an evolutionary increase in a certain parameter would not be simply the mirror image of a decrease in the same parameter. The asymmetry was thus passed upwards from the shortest time scale to the longest.

## 11 Conclusion

By considering organisms as dynamical systems, we are able to understand why they have the properties that Waddington portrayed in the epigenetic landscape. We can also make a number of inferences about development and evolution. Both large changes and multiple speciation should occur. Large changes, even more than small ones, are likely to be capable of being initiated by a number of different mutations. This greatly increases the likelihood that they will occur, certainly by comparison with regulatory gene mutations, which presumably involve one specific gene. Where a transition between two states is possible, it will be more likely to occur in one direction than in the other. In particular, phenocopying should occur in one direction only. This lack of symmetry will contribute to the irreversibility of evolution.

In biology, organisms are traditionally arranged into a hierarchical classification. A butterfly has certain properties, say its colour and markings, because of the species and variety it belongs to, whereas other properties, such as being segmented and having three

pairs of legs, are common to all insects. Still others, like the triplet code a butterfly uses to translate its genetic information, are properties of almost all organisms. We may feel we understand how DNA is transcribed and translated in a butterfly, but very little of our knowledge of this process was gained from work on butterflies.

Conversely, we often study a particular organism more because of the light it can throw on general problems than for its own sake. The aim of the vast amount of research that has been done on *Drosophila* has been to learn about genetics and development, not to satisfy an apparently limitless curiosity about fruitflies.

Now an organism is also a dynamical system. Hence even if biologists do not generally include dynamical systems as a level of the Linnaean hierarchy, it is still an appropriate level for studying biological phenomena, especially the most general ones that are shared by almost all organisms – and also some systems that we do not usually think of as organisms, like the Earth [5]. And, just as biologists do, we can also use particular dynamical systems as examples to help us in our work. We may think of equations (1) and (4) as mathematical fruitflies. We are interested in them not because they model important chemical reactions – it may even be that they do not model any real reactions at all – but because they are convenient to work with and yet share important properties with large classes of dynamical systems which almost certainly include many of those that occur in development.

Waddington saw the value of this approach half a century ago, which is why he devised the epigenetic landscape. Since then, mathematics has provided some of the techniques needed to make his idea work. We know far more about the behaviour of classes of dynamics, as distinct from individual equations [9]. We are also better able to judge which equations are typical of large classes and which are not. It may be a new departure that mathematics can play such a fundamental role in biology, but then it is a rather new kind of mathematics that is involved.

**Acknowledgement.** I am grateful to Mae-Wan Ho for helpful comments and for drawing Figure 1.

## References

- [1] Darwin, C. *The Origin of Species* (6th edition). London: John Murray, 1875.
- [2] Haldane, J.B.S. A Mathematical Theory of Natural and Artificial Selection. Part V. Selection and Mutation. *Proceedings of the Cambridge Philosophical Society* **23** (1927): 838-844.
- [3] Ho, M.W., E. Bolton and P.T. Saunders. Bithorax Phenocopy and Pattern Formation I. Spatiotemporal Characteristics of the Phenocopy Response. *Experimental Cell Biology* **51** (1983): 282-290.
- [4] Lewis, J., J.M. Slack and L. Wolpert. Thresholds in Development. *Journal of Theoretical Biology* **65** (1977): 579-590.
- [5] Lovelock, J.E. *The Ages of Gaia*. Oxford: Oxford University Press, 1988.
- [6] Mivart, St.G. *On the Genesis of Species*. London: Macmillan, 1871.

- [7] Peitgen, H.O., and P.H. Richter. *The Beauty of Fractals: Images of Complex Dynamical Systems*. Berlin: Springer, 1986.
- [8] Saunders, P.T. *An Introduction to Catastrophe Theory*. Cambridge: Cambridge University Press, 1980.
- [9] Saunders, P.T. Mathematics, Structuralism and the Formal Cause in Biology. In *Dynamic Structures in Biology*, edited by B.C. Goodwin, G.C. Webster and A. Sibatani, 107-120. Edinburgh: Edinburgh University Press, 1989.
- [10] Saunders, P.T. and M.W. Ho. Primary and Secondary Waves in Prepattern Formation. *Journal of Theoretical Biology* **114** (1985): 491-504.
- [11] Saunders, P.T. and C. Kubal. Bifurcations and the Epigenetic Landscape. In *Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems*, edited by B.C. Goodwin and P.T. Saunders, 16-30. Edinburgh: Edinburgh University Press, 1989.
- [12] Stewart, I. *Does God Play Dice? The Mathematics of Chaos*. Oxford: Blackwell, 1989.
- [13] Thom, R. An Inventory of Waddington Concepts. In *Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems*, edited by B.C. Goodwin and P.T. Saunders, 1-7. Edinburgh: Edinburgh University Press, 1989.
- [14] Thompson, J.M.T. and G.W. Hunt. *A General Theory of Elastic Stability*. London: Wiley, 1973.
- [15] Waddington, C.H. *Organisers and Genes*. Cambridge: Cambridge University Press, 1940.
- [16] Waddington, C.H. Genetic Assimilation of the Bithorax Phenocopy. *Evolution* **10** (1956): 1-13.
- [17] Waddington, C.H. *The Strategy of the Genes*. London: George Allen and Unwin, 1957.