

Spots and stripes: the evolution of repetition in visual signal form

Benjamin Kenward^{a,*}, Carl-Adam Wachtmeister^{a,1}, Stefano Ghirlanda^{b,1},
Magnus Enquist^{a,1}

^aDepartment of Zoology, Stockholm University, 106 91 Stockholm, Sweden

^bDepartment of Psychology, University of Bologna, Viale Carlo Berti-Pichat 5, 40127 Bologna, Italy

Received 23 May 2003; received in revised form 8 June 2004; accepted 14 June 2004

Abstract

It is common to find spatially repetitive patterns in animal visual signals. The evolution of such patterns is not well explained by existing theories of signal evolution. In this paper, we suggest that the evolution of signals with spatial repetition may be due to specific recognition problems and receiver biases. The logics of our hypotheses are studied in co-evolutionary simulations using artificial neural networks as models of receivers. These simulations yield repetitive visual signals under the following conditions: translations and reflections of the signal, partial obstruction of the signal, a fixed feature in the signal, and lateral inhibition in the receiver. In addition to regular repetitions our simulations sometimes result in other organisations of the signal such as blocky patterns and gradients.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Signal evolution; Repetition; Neural network; Simulation

1. Introduction

Spatial repetition is a common property of visual signals found in nature. Such signals may repeat details such as coloured stripes, dots or morphological structures. Signals with spatial repetition occur in most taxonomic groups that rely on visual communication, such as insects, slugs, fish, reptiles, birds and mammals (Cott, 1940 or see e.g. Coborn, 1991 (snakes); Wells and Clayton, 1993 (slugs); Lieske and Myers, 1994 (fish)). They also occur in some flowers and fruit. We define a repetitive pattern as one composed of a reoccurring sub-pattern, referred to as the generator (see Fig. 1). The goal of this paper is to assess the importance of a

number of factors for the evolution of spatial repetition in biological signals.

The evolution of signal form has recently received considerable attention (e.g. Endler, 1978, 1993; Guilford and Dawkins, 1991; Andersson, 1994; Bradbury and Vehrencamp, 1998; Endler and Basolo, 1998; Enquist and Arak, 1998; Boughman, 2002). Signal properties such as exaggeration, symmetry and colour have all been investigated but less focus has been placed on spatial repetition (but see e.g. Rothschild, 1964; Guilford, 1990; Guilford and Dawkins, 1991). It is well established that repetition in time, such as the displaying of the same signal over and over again, promotes detection (e.g. Haber, 1965; Haber and Hershenson, 1965; Bradbury and Vehrencamp, 1998). However, this is not always true for spatial repetition, which can sometimes in fact have the opposite effect, demonstrated for example by Cott's Zebra (Fig. 2) (Cott, 1940, but see also Windecker, 1939 in Järvi et al., 1981; Godfrey et al., 1987; Ruxton, 2002).

Factors influencing signal form can generally be grouped into two main categories. Firstly, a signal

*Corresponding author. Department of Zoology, University of Oxford, Oxford OX1 3PS, UK. Tel.: +44-1865-271245.

E-mail address: benjamin.kenward@zoology.oxford.ac.uk (B. Kenward).

¹InterCult, Department of History, Stockholm University, Kräftriket Hus 7B, 106 91 Stockholm, Sweden.




Fig. 1. Some example patterns that demonstrate the concept of repetition. Pattern a is repetitive with a generator of two units. Pattern b is technically repetitive, but generally referred to as a uniform pattern. We do not regard these patterns as interesting from the point of view of repetition. Pattern c has a transition from black to white, but is not repetitive as it has no repeated generator. Patterns d, e and f are additional examples of repetitive patterns. The generator of each repetitive pattern is underlined.




Fig. 2. Effect of repetition on signal detectability (based on Cott, 1940). Viewing this figure from increasing distances, the reader will notice that the repeated pattern fades into the background sooner than the other two.

may take part of its form from factors unrelated to communication. Examples are physical laws and constraints deriving from body structure or developmental programs. For example, with segmented animals such as caterpillars and earthworms, a signal that is produced at one segment may also appear at other segments, resulting in repetition. Even in unsegmented animals, reaction diffusion processes in development may result in repeated patterns (Turing, 1952). Also, features such as eyes which may comprise a part of a signal are limited in the extent to which their appearance can change, and this could influence the evolution of the rest of the signal.

Secondly, a signal's form will evolve to make it effective for communication and sometimes also to avoid detection by "non-intended" receivers (e.g. End-

ler, 1978, 1992; Bradbury and Vehrencamp, 1998). Several factors relating to communication can influence signal form. The signal must be detectable, transmitting effectively through an environment which tends to distort or attenuate the signal (e.g. Morton, 1975; Endler, 1987, 1993). The receiver can also influence signal form. Different receiver responses to variants of a signal act as selection pressures on signals. These differences in response may be adaptive (e.g. Zahavi, 1975, 1977; Andersson, 1994; Møller and Swaddle, 1997), or they may be due to non-adaptive receiver bias and a consequence of sense organs and neural architecture and neural processing (e.g. Staddon, 1975; Ryan, 1998; Enquist and Arak, 1998). Variation in signal images comes not only from variation in the population of signallers but also because one particular individual's signal can be perceived in numerous different ways, depending on conditions such as the distance between and the relative orientations of the sender and receiver (Endler, 1978; Enquist and Arak, 1998). Selection is based on the receiver's responses to all the different images of a signal it perceives. An example of how this can influence signal form is the promotion of symmetry when a signal must be recognized when projected in different orientations on the retina (Enquist and Arak, 1994).

We consider here a number of very general communication related hypotheses for the evolution of repetition, and investigate their logic using evolutionary simulations. The factors considered include discrimination tasks, the effect of obstruction of the signal, various transformations such as translations and reflections, lateral inhibition and receiver biases. We use artificial neural networks as models of receiver mechanisms (Ghirlanda and Enquist, 1998; Kamo et al., 1998; Phelps, 2001) in simulations of coevolutionary processes between senders and receivers. Receiver behaviour is treated as innate and modified by evolution. In addition to spatial repetition, we also consider some other patterns that emerged in our simulations and are also found in nature, including symmetries, blocky patterns and gradients.

2. Simulation model

2.1. The neural network model

We use feed-forward neural networks to simulate signal receivers making decisions based on visual stimuli. These models are inspired by the parallel structure of real nervous systems and have been used successfully to reproduce empirical results from studies of stimulus control in real animals, including such effects as generalization, peak-shift (super-normality) and stimulus intensity effects (Ghirlanda and Enquist,

1998; Kamo et al., 1998; Phelps, 2001). Because the models' purpose is to reproduce behaviour, parameters are set to reproduce stimulus response results, rather than to match physiology. The basic idea behind this approach is therefore to design a model which replicates empirical results, and then investigate how the model responds under a variety of simulated conditions.

The stimulus is received in the first layer of the network, which models a retina. The retinal units then stimulate the cells in the next layer, known as the hidden layer, through weighted connections. The hidden layer cells are connected to a single output cell, whose value is regarded as the network's response to the stimulus. It is the values of the connection weights which determine the network's response to a particular stimulus. Our networks have one hidden layer, and are identical to those used in a number of other studies (for details see e.g. Enquist and Arak, 1998). Unless otherwise stated, we use the following parameters: Networks have either 24 retinal units and 10 hidden layer units, or 12 and 6. There is one output unit, which can produce a response between 0 and 1. The signal is composed of 12 units. The retina and signal are one dimensional, and their units' values are bounded between 0.01 and 1. In the following text we refer to values of 1.0 as white, and values of 0.01 as black.

We also examine the effects of using networks that have inbuilt lateral inhibition (for a review of the phenomena see e.g. Coren et al., 1999). In these simulations the signal stimulates the input layer of the network only after having been subject to a lateral inhibition algorithm. The algorithm reduces the strength of a signal unit proportionally according to the intensity of its two neighbouring cells. The strongest possible perceived intensity comes from a cell receiving intensity 1, with neighbours both receiving 0.01. The lowest possible perceived intensity comes from a cell receiving intensity 0.01, with neighbours both receiving 1. The strength of inhibition is governed by a coefficient a . If S_i is a unit in the image and L_i is the corresponding unit after lateral inhibition, then:



$$L_i = (1 - a)S_i + a\left(\frac{1}{2} + \frac{S_i - S_{i-1}}{4} + \frac{S_i - S_{i+1}}{4}\right). \quad (1)$$

Meaningful values of a are between 0 and 1. A value of 0 means no lateral inhibition. A value of 1 is so much that a row of black units and a row of white units are both perceived as 0.5 strength by the network.

2.2. Stimuli

In most of the simulations the task of the network is to discriminate a signal (S+), from both a background (B) and another stimulus (S-). The S+ and S- vary from simulation to simulation but the background is always uniformly 0.01 (black) and covers the whole



Fig. 3. An example pattern. The generator  occurs 4 times, meaning 8 out of 12 units are covered by a generator, so the repetition count is 8. Other generators are also in the pattern, such as , which only covers 6 units.

retina when it is presented. The S- is a uniform pattern of value 1 (white), and the same size as the S+. From each stimulus one or more retinal images are created, by presenting the stimulus to the retina in different ways (see below).

To measure repetition, we calculate how much of a pattern is made up by a generator (remember a generator is a repeated sub-pattern). If there are several generators, the one covering most of the pattern is chosen. The number of units made up by a generator is referred to as the repetition count. For an example, see Fig. 3. A perfectly repetitive signal of 12 units will have a repetition count of 12, as the signal will be composed only of the generator. Although patterns which are completely uniform could be regarded as repetitive (with a uniform generator), we assign these patterns a repetition count of zero. Most patterns that emerge in the simulations contained only black and white units. Therefore, to simplify the identification of generators, grey units are treated as black or white, depending on whether they are below or above 0.5.

2.3. Fitness

To behave optimally, the receiver must respond with an output of 0.9 to each of the images produced by presentation of the signal S+ on the retina (see below). Let us write as e_i the absolute value of the difference between the optimal output and the actual output for image i . As a measure of receiver performance we first calculate the geometric mean of $1 - e_i$:

$$f_{S+} = \sqrt[n]{\prod_{i=1}^n (1 - e_i)}, \quad (2)$$

where n is the number of images produced by the S+. The geometric mean is used because, as in real life, this leads to severe fitness consequences for a highly inappropriate response to even one of the images. According to Eq. (2) f_{S+} is equal to 1 if the response to every image is perfect (all e_i are zero), and decreases as any of the errors e_i increase. Receivers should also respond with an output of zero to all the images produced by the stimulus S- and the background B. We can then calculate f_{S-} and f_B as above. We combine these three performance measures to obtain receiver fitness like this:

$$f = f_{S+}(f_{S-}f_B)^5. \quad (3)$$

f_{S-} and f_B are raised to the 5th power, meaning the reactions to the background and S– make a greater contribution to receiver fitness than the reaction to the S+, because in nature stimuli to be reacted to (predators, food items, mates) are typically encountered less often than stimuli which should be ignored, such as the background. The precise value of the exponent only influences the speed of the evolutionary process and not its outcome (so long as it is greater than 1).

The fitness of the sender is also based on the receiver responses, and it is also calculated with Eq. (2), with the difference that the optimal output is 1 rather than 0.9. This means there is a small amount of conflict between the two players—the sender wants to elicit the maximum possible response, but the receiver wants to respond slightly below maximum. We include a small amount of conflict because it is biologically realistic, and because in previous studies we had discovered that simulations with conflict are less likely to get stuck in a configuration where the receiver is unable to distinguish the signal from the negative stimuli. The value of 0.9 is to some degree arbitrary but chosen as a trade off between the need to include conflict, and the need to reward discrimination, i.e. the response to the signal being maximally different from the response to the background. Pilot experiments revealed that the model is not sensitive to small changes in this value.

2.4. Iterative evolution

To train a naïve network to respond appropriately to signals, we use a process known as iterative evolution, which can evolve a network that solves a discrimination problem and a signal that efficiently stimulates a network, in a coevolutionary process. For an example of previous use of this process, see [Enquist and Arak \(1998\)](#). Network and signal evolution both work in the same way—a new individual is produced each generation by mutating the existing one. Both individuals are tested and allocated fitnesses according to the rules of the particular simulation. If the new mutant is more fit, it is kept and the old individual is discarded. If the mutant is inferior, it is discarded and a new mutant will be generated next generation. Most of our simulations involve coevolution between signal and receiver. The receiver undergoes the process of mutation, testing, and possible replacement first. The signal then undergoes the same process.

Every generation each network connection has a probability of mutating of 3 divided by the number of connections, meaning that on average there will be 3 mutations each generation. The changes in connection weight are taken from a normal distribution with mean 0 and standard deviation 0.05. Every generation each signal unit has a probability of mutating of 1 divided by the number of signal units, meaning that on average

there will be 1 mutation each generation. The changes in unit value are taken from a normal distribution with mean 0 and standard deviation 0.05. These particular parameters were chosen because experience showed that they create a balanced coevolutionary situation in which neither the sender nor receiver has an advantage of evolving faster than the other.

2.5. Coevolutionary scenarios

Each coevolutionary simulation was run for 50,000 generations, which was long enough for the signal to reach a relatively stable form (fitnesses changed by less than 0.01 in the last 1000 generations in 99% of all simulations). In some simulations functional signals did not evolve (indicated by low sender and receiver fitnesses). We discarded these simulations from further analysis, because we are only interested in functional signals. For each parameter set, we ran each simulation until we had evolved 250 non-excluded signals.

2.6. Statistical analyses

We needed a way to analyse the results of our simulations to determine if repetitive patterns occurred at a different frequency than expected by chance. To determine what would be expected by chance in random patterns, we generated every possible random pattern and calculated the frequency distributions of each repetition count. We also calculated these frequencies for the signals evolved from each parameter set, and then compared the observed frequencies with the expected frequencies using chi-squared tests. This method also allowed us to determine if non-repetitive patterns (those with low repetition counts) were evolving more often than expected by chance.

Several simulations resulted primarily in symmetrical signals, which by their nature have higher repetition counts than non-symmetrical ones. To test if these sets of signals had frequencies of repetition counts which could not be explained by their symmetry alone, we compared them with random patterns which were also symmetrical. To do this, we calculated the proportion of evolved signals which were perfectly symmetrical, and biased the distribution of random patterns by removing non-symmetrical patterns at random until it had the same proportion of symmetrical patterns as the set of evolved patterns did. The evolved repetition counts and random (but biased) repetition counts were then compared using chi-squared tests as before.

3. Hypotheses and simulations

We now describe our hypotheses, simulations of these hypotheses, and present and discuss the results. We

begin with an investigation of the effects of the need to discriminate signals from the background and other stimuli. Next, we consider how in nature a signal can be presented in more than one way, depending on orientation, distance, etc. This results in different images which require the same response. The way in which these images are produced can be described by transformations of various sorts (e.g. Stewart and Golubitsky, 1992). The transformations we consider are translation and reflection. We then investigate the effect of partially obstructing the signal. Next, we investigate some effects of pre-existing structure of both receiver and signaller—in the case of the signaller the existing structure is a fixed feature such as an eye; for the receiver the structure is the lateral inhibition effect common to most visual systems. Finally, we investigate whether repetition can be a superstimulus.

3.1. Hypothesis 1—need for signal uniqueness

One reason for the evolution of a non-uniform repetitive pattern could be the need for distinctiveness from already existing uniform patterns. Hurd et al. (1995, Fig. 3) showed that signals which must be distinguished from each other will evolve to be as different as possible. Consider, for example, two fixed signals, one of which is entirely black, and one of which is entirely white. A third signal, which evolves to be as different as possible, will often evolve to be black and

white, rather than grey. We suggest that because this need to be unique can produce signals with patches of different colour, this can lead to repetition.

3.1.1. Simulation method

We use the standard parameters of a signal size 12, retina size 12, a black background, a fixed white S−, and an S+ (the signal) which coevolves with the receiver network.




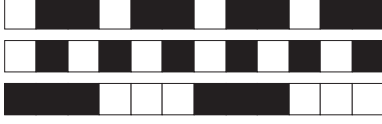
3.1.2. Results and discussion

Table 1 shows a few signals which result from this simulation. It can be seen that black and white units alternate, which can create somewhat repetitive looking patterns. However, the distribution of repetition counts obtained is not significantly different than expected by chance in random black and white patterns ($\chi^2_6 = 6.66$, NS). With 12 signal units about 2% of patterns will be perfectly repetitive if black and white units are randomly assigned to the pattern.

3.2. Hypothesis 2—translating the signal across the retina

Objects do not always project the same unvarying image on the retina. An object will project in different positions depending on the relative positions of the object and observer. Even in animals which direct the central part of the retina (the fovea) towards objects of

Table 1
Typical signals obtained from simulations 1 and 2

Effect simulated	Description of patterns obtained	Example of patterns	Repetition count	Occurrence (% of runs)
1—Pressure for uniqueness	Random black and white		Varies randomly	100
	Blocky		0	33.6**
2—Translation	Almost perfectly repetitive		10	18.8*
	Perfectly repetitive		12	7.2**

N is 250 for each simulation.
*Indicates a significantly high occurrence of $p < 0.05$.
**Indicates $p < 0.001$. For details of statistics see text.

interest, a signal image is unlikely to be projected so that each photoreceptor always receives exactly the same input as with a previous image. With a repetitive pattern, any small part of the retina which is exposed to a portion of the signal will be stimulated in a similar way, regardless of the position of the signal. This is known as invariance of the signal. This may favour the evolution of repetitive patterns because they are effective at stimulating the receiver regardless of the position of the signal.

3.2.1. Simulation method

We increase the retina size from 12 to 24, and present the signal in all possible positions on the retina to create the signal images, resulting in 12 different images. The 12 units of each signal image not comprised of the signal itself are set to the same value as the black background. Twelve white S– images are created in the same way. We refer to this treatment as translating the signal across the retina. Note that in this simulation, and in all other coevolutionary simulations, the S– and background are always present, meaning there is always a pressure to be unique.

3.2.2. Results and discussion

The resulting signals are again mainly black and white, with very few grey units (see Table 1). Some of the signals are random patterns, as in the previous simulation. But a large number seem to have identifiable structure. These structured signals fall into two categories. The first comprises signals organized into two or three areas of black and white, which we call ‘blocky’. These signals have a repetition count of zero, which occurs significantly more often than expected by chance ($\chi^2_1 = 726.65$, $p < 0.001$). They can be discriminated by the receiver from the S– because the white blocks are smaller, and because the signals sometimes contain more than one white block.

The second category comprises repetitive signals (see Table 1). Perfectly repetitive signals (repetition count 12) appear significantly often ($\chi^2_1 = 40.99$, $p < 0.001$), as do nearly perfectly repetitive signals (repetition count 10) ($\chi^2_1 = 5.33$, $p < 0.05$). We believe the cause of this repetition could be what we call ‘self-mimicry’ within the signal. Imagine that there is a sub-pattern within the signal which is more effective than any other sub-patterns at stimulating certain parts of the retina. The signal will become more effective if other regions of the signal evolve to ‘mimic’ this efficient sub-pattern. Due to response biases, there may be possible sub-patterns which are even more efficient, but these are likely to be similar to the currently existing best sub-pattern. As this happens, the receiver will adapt to better recognize the particularly efficient sub-pattern. These two processes will reinforce each other in co-evolution, meaning that a particular sub-pattern can become repeated throughout

the signal. If during evolution a new sub-pattern arises which is even more efficient than the current most efficient one, then that pattern will then be mimicked instead.

3.2.3. Analysis of sensitivity of the model to the properties of the S–

We ran an additional simulation identical to the one just described, but with an additional S– composed of random units from a uniform distribution. Visual inspection of the signals obtained showed there was little or no difference from simulation 2, and the distribution of repetition counts obtained was not different ($\chi^2_6 = 4.607$, NS).

3.3. Hypothesis 3—reflection of the signal

When signal images are created on a retina, as well as varying in position, they can also vary in orientation. For example, a fish might swim past from right to left, or left to right, creating images which are reflections of each other. Repetitive signals can be invariant when reflected (symmetrical), so reflection may favour the evolution of repetitive patterns.


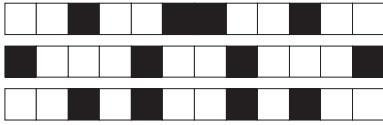
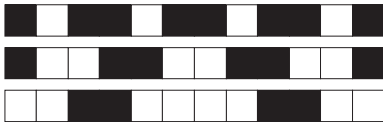
3.3.1. Simulation method

There are two variants of reflection used. The first reflection simulation (3a) is the same as simulation 1 (need for uniqueness), but there are two images of the signal, the second created by reflecting the first. The second reflection simulation (3b) is the same as simulation 2 (translating the signal) but combines reflection with the translation, meaning there are 24 images, 12 of them reflected versions.

3.3.2. Results and discussion

Simulation 3a results in signals which are all symmetrical, i.e. the signal and its mirror image are identical, (see Table 2). We explain this as follows. Whenever a modification in the signal appears (because of the pressure to be unique) it would increase the fitness of the signal further to incorporate the mirror modification on the other half of the signal (see also [Enquist and Arak, 1998](#), Fig. 2.17). Perfectly repetitive signals and nearly perfectly repetitive signals appear more often than expected by chance ($\chi^2_1 = 58.23$, $p < 0.001$ and $\chi^2_1 = 223.47$, $p < 0.001$, respectively). There were also a higher number of blocky signals than expected by chance ($\chi^2_1 = 104.99$, $p < 0.001$). However, it is possible that both these effects are due to the inherent properties of symmetrical signals. We control for this by biasing the expected by chance distribution to have the same amount of symmetry as the results we obtained. We then compare this distribution with the obtained results as before. There is no significant difference ($\chi^2_6 = 11.61$, NS). This shows that the increased numbers of both

Table 2
Typical signals obtained from simulations 3a and 3b

Effect simulated	Description of patterns obtained	Example of patterns	Repetition count	Occurrence (% of runs)
3a—Reflection and 3b—Reflection with translation	Symmetrical, blocky		0	3a: 14.8**
		3b: 39.2**		
	Symmetrical, almost perfectly repetitive		10	3a: 30.4**
		3b: 26.4**		
	Symmetrical, perfectly repetitive		12	3a: 14.4**
		3b: 16.0**		

N is 250 for each simulation.

**Indicates a significantly high occurrence of $p < 0.001$. For details of statistics see text.

repetitive and blocky signals is indeed the result of the inherent properties of symmetrical patterns, and nothing more.

Simulation 3b (translation with reflection), results in signals which are similar to those obtained from translation without reflection (simulation 2), with the difference that the majority are symmetrical (see Table 2). There are some with no obvious structure besides symmetry, but most have additional structure. Many more are blocky than expected by chance ($\chi^2_1 = 1019.77$, $p < 0.001$). Many are perfectly or nearly perfectly repetitive ($\chi^2_1 = 283.86$, $p < 0.001$ and $\chi^2_1 = 33.59$, $p < 0.001$, respectively). In fact this simulation produces significantly more perfectly repetitive signals than translation without reflection ($\chi^2_1 = 28.97$, $p < 0.001$). As we saw in the previous reflection simulation, symmetry alone can produce both of these effects. But this time, after controlling for symmetry, there are still a significantly large number of blocky patterns ($\chi^2_1 = 128.92$, $p < 0.001$) and perfectly repetitive patterns ($\chi^2_1 = 5.87$, $p < 0.05$). It seems reasonable to assume that these signals are produced by the same mechanisms as when there is no reflection. The blocky patterns are distinguishable from the S– because of their smaller white areas. The repetitive patterns are probably the result of self-mimicry, enhanced by the symmetry.

3.4. Hypothesis 4—partial obstruction of a signal

A problem facing signal senders is to be recognized when parts of the signal are obscured. For example, a

snake may need to be recognized although much of it is covered by vegetation. This need could be responsible for repetitive patterns, which look similar even when parts are obscured.



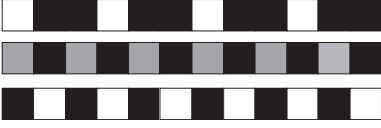



3.4.1. Simulation method

The simulation is run with retina, signal, background (black) and S– (white) all of 12 units. Versions of the signal and S– are created by obscuring a number of units between 0 and 8, in all possible ways, but with the restriction that the visible part of the stimulus was never divided into more than one area. The obscured units are set to the black background colour. Each obscured version of the stimulus is then translated across the retina, as before, to create the images. The smallest amount of signal visible in an image is therefore 4 units.

3.4.2. Results and discussion

Almost all the resulting signals (see Table 3) fall clearly into one of three categories. Those in the first category are a uniform grey colour, almost always between 0.1 and 0.3. As explained earlier, black and white patterns are more different from uniform black and uniform white than uniform grey is. However, now that the receiver is presented with only parts of the signal, it seems that the grey is effective enough to evolve, presumably because any part of the grey signal is recognizably grey. It is not so simple to find a black and white pattern which is easily recognizable from any part of it.

Table 3
Typical signals obtained from simulations 4 and 5

Effect simulated	Description of patterns obtained	Examples of patterns	Repetition count	Occurrence (% of runs)
4—Partial obstruction	Uniform grey		0	40.0
	Gradient		0	37.0
	Perfectly repetitive		12	13.6**
5—Fixed feature	Blocky		4	17.2**
	Almost perfectly repetitive		10	24.8**
	Perfectly repetitive		12	13.2**

N is 250 for each simulation.

**Indicates a significantly high occurrence of $p < 0.001$. Significances cannot be calculated for the grey and gradient patterns, because in these simulation more than one type of pattern arises with repetition count 0. For details of statistics see text.

The next category of patterns, which do not occur so often, show a gradient across the signal. This gradient is usually very shallow (varying only 0.05 from one end to the other) though it can be steep, ranging almost from black to white. This signal form is probably a viable solution because any smaller part of it is still a gradient, distinguishing it from the background and S-. The grey and the gradient signals both have repetition counts of zero, because they are not repetitive. The high occurrence of zero repetition count is significant ($\chi^2_1 = 4506.57$, $p < 0.001$), though perhaps not particularly informative because it covers two different categories of signal.


The third category is of perfectly repetitive signals. More perfectly repetitive signals result from obstruction, than from translation alone (simulation 2) ($\chi^2_1 = 15.33$, $p < 0.001$). This could be due to the fact that each small part of the signal is now subject to the pressure to be unique from the background and S-. So black and white must be spread fairly evenly throughout the signal to ensure that each possible small part contains both black and white. However, there are numerous possible signals meeting this criterion which are much less repetitive than those which evolved. There must, therefore, be another force driving the signals towards

repetition. This, we believe, is the self-mimicry force, enhanced by the obstruction, because when a sub-pattern arises which is particularly effective at stimulating the receiver, it is only present in some of the possible obscured images. There is now a very strong selection pressure on the signaller to be able to stimulate the retina with all parts of the signal. The easiest way for this to occur is if the inferior sub-patterns are replaced by the effective one, meaning that the retina does not have to solve the difficult task of recognizing a large number of different images.

3.5. Hypothesis 5—a fixed feature in part of the signal

A signal might contain within it a fixed feature, which has a function unrelated to signalling, but is recognizable. For example, an eye on the side of a fish will be minimally affected by evolutionary pressures on signal form, despite the fact that the whole shape and colour of the fish is a signal. The eye must be able to see, so the range of forms it can take is severely limited. An evolutionary stable property of a signal might be used by the receiver for recognition, and the effectiveness of the signal might therefore be increased if the form of the eye is mimicked elsewhere in the signal.

3.5.1. Simulation method

The simulation is identical to translation alone (simulation 2), with the exception that the rightmost three units of the signal are locked as . The remaining nine units of the signal are allowed to evolve as normal.

3.5.2. Results and discussion

The signals produced (see Table 3) are the same as those resulting from translation alone (simulation 2), that is, blocky signals and repetitive signals, as well as some random patterns. The blocky signals in this case have a repetition count of four, not zero, because of the fixed feature, and they appear significantly often ($\chi^2_1 = 71.18$, $p < 0.001$). More perfectly repetitive patterns are produced in this simulation than in translation alone ($\chi^2_1 = 25.14$, $p < 0.001$). This supports the hypothesis that a fixed feature can produce repetition.

3.6. Hypothesis 6—networks with lateral inhibition

An effect which could be important in the evolution of repetitive patterns is the physiological process of lateral inhibition. This occurs when retinal cells which are stimulated inhibit neighbouring retinal cells. The more a cell is stimulated, and the closer it is to another cell, the more intensely it will inhibit the other. This leads to retinal cells near boundaries between light and dark areas having greater response intensity, because there are not so many strongly stimulated cells in their neighbourhood (e.g. Coren et al., 1999). It is suggested that by enhancing edges, this process makes it easier to distinguish objects from backgrounds in varying light conditions (e.g. Bruce et al., 1996). This might also lead to networks being highly stimulated by a pattern containing many transitions between light and dark, such as repetitive patterns. Stripes, therefore, have a greater apparent maximum intensity than solid patterns. Guilford (1990) points out that this intensity could be responsible for the results of Schuler and Hesse (1985), who discovered that naïve chicks (*Gallus gallus domesticus*) had a greater hesitancy to eat black and yellow striped prey than green prey.

3.6.1. Simulation method

We modify some of the previously described simulations by using networks which have inbuilt lateral inhibition (as described in the method section above). The first lateral inhibition simulation (6a) has the same parameters as the one used to demonstrate the effect of uniqueness (simulation 1). The retina, signal, and S– are all of 12 units. The lateral inhibition coefficient, a , is 0.6. The second simulation (6b) combines the effects of lateral inhibition and translation. The parameters are the same as in simulation 2—the retina is of 24 units, signal and S– of 12. We also investigate the effect of the

strength of lateral inhibition by using different values of a of 0.2, 0.4, and 0.6 (Greater values of a give stronger lateral inhibition).

3.6.2. Results and discussion


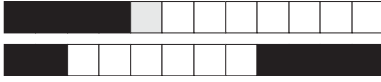
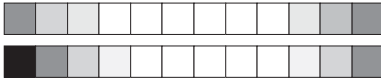


Simulation 6a, which has lateral inhibition but no translation of the signal, does not result in a distribution of repetition counts significantly different to that expected by chance ($\chi^2_6 = 9.11$, NS). Even strong lateral inhibition without translation is therefore not enough to produce repetition (see Table 4). The results of adding translation to the lateral inhibition in 6b depended on the value of a . A value of 0.2 produces signals which are not significantly different from those produced with translation alone ($\chi^2_6 = 10.39$, NS). Values of a of 0.4 and 0.6, however, produce significantly more perfectly repetitive signals than without lateral inhibition ($\chi^2_1 = 19.40$, $p < 0.001$ and $\chi^2_1 = 121.23$, $p < 0.001$, respectively). This suggests that the more lateral inhibition there is in a network, the stronger the selection for repetitive patterns.

Other signals are also produced. Blocky patterns and gradients are both present. The blocky patterns are most common when lateral inhibition is weak (i.e. a is small) ($\chi^2_2 = 66.5$, $p < 0.001$). This is consistent with the hypothesis that stronger lateral inhibition should produce more transitions between colours. Most of the gradient patterns are different from the gradients found in other simulations, because they shift from dark at the edge to white in the middle and then back to dark again, rather than being a simple and consistent gradient from one end to the other. They are most common when lateral inhibition is strong ($\chi^2_2 = 12.6$, $p < 0.01$). The gradient signals and the white S– are different in the following way, which presumably allows the receiver to distinguish them: The gradient signals, though white in the middle, have no sharp transition between white and black at the edges because they fade to black. The S– images, on the other hand, always contain a sharp transition between the white edge of the signal and the black background. Due to lateral inhibition, the edges of the S– images therefore stimulate the receiver much more intensely than edges of the gradient signal images. The signal is still distinguishable from the black background because it does contain white units. This explanation of the result implies the receiver has evolved to produce a high output in response to a medium input and a low output in recent to a high input—to what extent this is biologically realistic is uncertain.

3.7. Hypothesis 7—receiver bias for an increased number of components

Observations from experimental psychology (Hanson, 1959; Mackintosh, 1974) and ethology (Tinbergen, 1951; Hinde, 1970) have taught us that some unfamiliar

Table 4
Typical signals obtained from simulations 6a and 6b

Effect simulated	Description of patterns obtained	Examples of patterns	Repetition count	Occurrence (% of runs)
6a—Lateral inhibition, $a = 0.6$	Random black and white		Varies randomly	100
6b—Lateral inhibition with translation	Blocky		0	$a = 0.2$: 24.8 $a = 0.4$: 11.6 $a = 0.6$: 0.8
	Gradient		0	$a = 0.2$: 13.2 $a = 0.4$: 22.0 $a = 0.6$: 25.6
	Almost perfectly repetitive		10	$a = 0.2$: 20.4* $a = 0.4$: 14.8 $a = 0.6$: 31.6**
	Perfectly repetitive		12	$a = 0.2$: 7.2** $a = 0.4$: 14.4** $a = 0.6$: 25.2**

N is 250 for each simulation.



*Indicates a significantly high occurrence of $p < 0.05$.

**Indicates $p < 0.001$. Significances cannot be calculated for the blocky and gradient patterns, because in these simulation more than one pattern arises with repetition count 0. For details of statistics see text.

stimuli can elicit stronger responses than familiar ones. If a particular pattern were a stimulus to an animal, then repetition of that pattern could be a superstimulus. For example, [ten Cate and Bateson \(1989\)](#) showed that male Japanese quail chicks (*Coturnix coturnix japonica*) that had been imprinted on white adults with artificial dots painted on tended to prefer adults which had even more dots than the familiar birds. Receiver bias for increased numbers of signal components could therefore promote the evolution of more repetitive signals.

3.7.1. Simulation method

We demonstrate a similar result here with two simulations. We are now simulating a learning situation, in which the signal does not evolve. Instead, we train the network to react to an unchanging stimulus, and then analyse the network's responses to some unfamiliar stimuli. In the first model, the retina is of 12 units. There is an S+ which is one white unit. There are 12 images of this stimulus, produced by translating it across the

retina. The network's task is to distinguish the S+ from the black background by reacting maximally to the stimulus and not at all to the background. There is no S− in this simulation, nor is there coevolution—only the network evolves. Once the network's fitness has reached 0.5, evolution is stopped and the network is tested to see how strongly it reacts to a number of novel stimuli with more than one white unit. The stimuli tested are repetitive patterns and uniform patterns. In the second model we add an S− of two white units, like this:  (we also keep the black background). The S+ in this model is two white units separated by one black: .

3.7.2. Results and discussion

In the first model, it was found that the network reacted more strongly to patterns with more white units ([Fig. 4a](#)). This is because adding more white units increases the difference between the background and the stimulus. However, increasing the number of white units in a uniform stimulus and increasing the number of

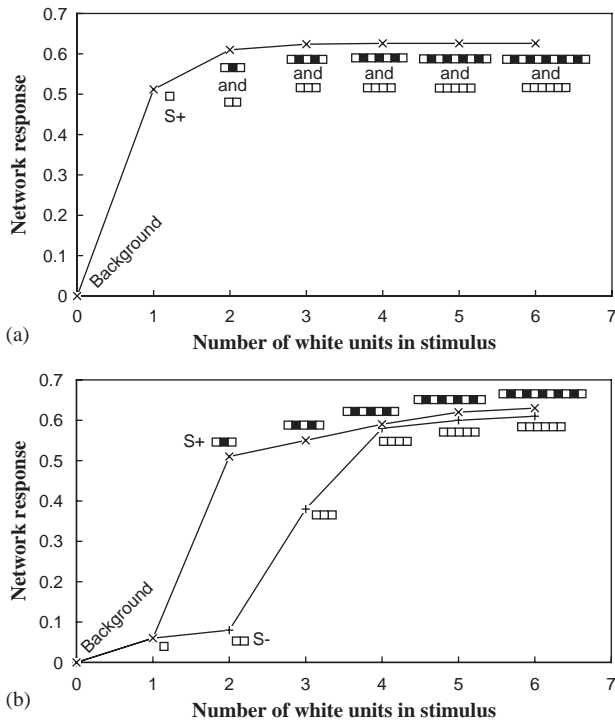


Fig. 4. The responses of a trained network to a number of test stimuli. (a) The network is trained to react to an S+ of one white unit, and not to react to a black background. The response to test stimuli increases with the number of white units, but is not dependent on whether the units are arranged in repetitive or uniform patterns. (b) This time the network is trained to react to an S+, but not to a black background or an S-. The S+ and S- both contain two white units, but in the S+ they are separated by a black unit. The response to test stimuli again increases with the number of white units, but this time there is also a preference for repetitive patterns over uniform ones.

white units in a repetitive stimulus has the same effect, demonstrating that a repetitive pattern is not the only pattern type that will super-stimulate this receiver. In the second model after training the bias for repeated white units becomes stronger than the bias for uniform white patterns (Fig. 4b). This demonstrates that if a signal includes a unit which is repeated just once, a bias can occur for signals which have more repetitions of that unit (again because they are more different from the background). The presence of an S- which is two adjacent white units means the bias for repetitions is greater than the bias for larger uniform patterns.

4. General discussion

4.1. Summarising the results

Using simulations with artificial neural networks we have demonstrated the logic of a number of hypotheses for factors that might be responsible for the evolution of spatially repetitive signals. All the factors which we simulated resulted in the increased probability of the

evolution of repetition, although no treatment produced solely repetitive patterns. We identified the following factors as relevant for the evolution of repetition:

1. The effect of the need for uniqueness. Due to the need to discriminate between different stimulations, response biases appear in the receiver, which can result in signals taking on unique patterns containing transitions between colours (see also Hurd et al., 1995). Patterns which contain transitions can by chance be repetitive. In nature, other forces such as developmental programs could operate to make transition containing patterns more likely to be repetitive.

2. The effect of translating the signal. When the signal is projected onto the retina in different positions, repetitive patterns evolve. This is probably because repetitive signals are invariant—that is they appear similar in different positions.

3. The symmetry effect. Due to the receiver's need to recognize both the signal and its mirror image, the receiver develops a bias for bilateral symmetries. This increases the amount of repetition because of the inherently repetitive properties of symmetrical signals.

4. The effect of partial presentation of the signal. When the image on the retina comprises only part of the signal, for example because it is obstructed, repetitive patterns evolve. This is probably because a repetitive pattern can be identified without needing to see the whole.

5. The fixed feature effect. If there is a morphological feature which cannot evolve due to constraints, then the rest of the signal might mimic that feature, producing repetition.

6. The lateral inhibition effect. Networks with lateral inhibition are more strongly stimulated by patterns containing transitions between colours, and one way for a pattern to contain many transitions is for it to be repetitive.

7. The effect of a preference for several signal components. A network, which is trained to prefer a stimulus composed of two separate components to a uniform stimulus, subsequently prefers stimuli with greater numbers of repeated components to uniform stimuli. A preference of this type could result in the evolution from a signal with a few repeated components to a signal with many repeated components.

We believe that some of the above effects, such as translation and partial obstruction, may be mediated by a more proximate evolutionary process which we call self-mimicry. There will be inherent biases in the receiver for particular parts of a signal. The signal will therefore become more efficient if it repeats any particularly effective subsection. Due to the receiver's need to recognize a signal presented in different positions on the retina, or only partially presented, the receiver will also benefit by specializing in recognizing any particular signal part which is repeated.

These results are consistent with a number of similar theoretical studies, based on simulations with neural networks, which demonstrate how preferences appearing as by-products of recognition problems affect the evolution of signal form (Enquist and Arak, 1994; Johnstone, 1994; Hurd et al., 1995; Bullock and Cliff, 1997; Enquist and Johnstone, 1997; Enquist and Arak, 1998; Kamo et al., 1998; Phelps, 2001).

4.2. The other common patterns—blocks and gradients

Though we began this investigation with the intention of shedding some light on how repetition evolves, we have also evolved other types of pattern, most noticeably patterns with blocks of colour, and patterns with gradients. These are all patterns which are often seen in nature, and the fact that they appear in our simulations lends legitimacy to our models. We have not investigated in detail the reasons why blocky and gradient patterns evolve, and further studies in this area may be worthwhile.

4.3. Additional hypotheses for the evolution of repetitive patterns

Apart from those causes that have been analysed above there are a number of others that could be responsible for the evolution of repetitive signals. Of course, not all repetitive patterns are signals. In fact repetition may more often serve to conceal than advertise (e.g. Thayer, 1918; Cott, 1940). Ortolani (1999) and Ruxton (2002) review a wide range of possible functions of repetitive patterns in carnivores and zebra, respectively. It is beyond the scope of this paper to comprehensively list all hypotheses related to communication but we state three more below.

Firstly, regular patterns, such as repetitive patterns, might be efficient since they are sometimes easier to detect in the natural and often chaotic stimulus environments—camouflage is often attained by disrupting regularity (e.g. Thayer, 1918; Cott, 1940). Second, a sender could exploit the fact that reality and perception sometimes disagree, referred to as visual illusions (e.g. Coren et al., 1999). For instance, a two-dimensional stripy pattern with horizontal bars will look more broad than the same sized pattern with vertical bars, (although depending on the density of bars, this effect can sometimes operate in the opposite direction, see e.g. Coren and Girgus, 1978). Third, a pattern can be highly conspicuous at short distance, but at the same time cryptic at longer distances (e.g. Cott, 1940; Rothschild, 1964; Endler, 1978; Järvi et al., 1981). Striped patterns are a good example; at close range they are very conspicuous, but at long range the stripes blur together so that, for example, a black and white striped area appears grey (e.g. Mottram, 1916; Cott, 1940; Endler,

1978; Coren et al., 1999). For example, Windecker (1939, in Järvi et al., 1981), showed that yellow and black striped cinnabar moth larvae (*Tyria jacobaea*) are cryptic when seen at distance amongst the flowers of their host plant (*Senecio jacobaea*), but are aposematic when seen close up. Aposematism is probably not the only situation in which crypsis at long distance and conspicuousness at short distance is advantageous. Imagine for instance a courting male that needs to display to females without attracting predators. Predators often have to recognize stimuli from a greater distance than the distance the prey (e.g. courting male) has to its signal target (e.g. Endler, 1978, p. 334).

4.4. Conclusions

Repetition is an important ingredient in signals which have often been analysed, such as the peacock's tail, but it is an ingredient that has not often been explained. There are a number of hypotheses for its evolution relating to different aspects of the communication process. Introducing a more realistic mechanism into coevolutionary simulations has given us insights into these hypotheses which are difficult to obtain using other tools such as game theory.

Acknowledgements

This research was supported by grants from Riddarhuset, Riksbanken Jubileumsfond and Marianne och Marcus Wallenbergs Stiftelse. Ben Kenward was supported by the Dumbleton Trust. We are grateful to Robert Kenward, Andy South, and several anonymous referees for comments on the manuscript.

References

- Andersson, M., 1994. Sexual Selection. Princeton University Press, Princeton.
- Boughman, J.W., 2002. How sensory drive can promote speciation. Trends Ecol. Evol. 17, 571–577.
- Bradbury, J.W., Vehrencamp, S.L., 1998. Principles of Animal Communication. Sinauer Associates Inc., Sunderland, MA, pp. 1–882.
- Bruce, V., Green, P.R., Georgeson, M.A., 1996. Visual Perception. Physiology, Psychology, and Ecology. Psychology Press, Hove.
- Bullock, S., Cliff, D., 1997. The role of “hidden preferences” in the artificial co-evolution of symmetrical signals. Proc. R. Soc. London Ser. B 264, 505–511.
- Coborn, J., 1991. Atlas of Snakes of the World. TFH Publications, Neptune, NJ.
- Coren, S., Girgus, J.S., 1978. Seeing is Deceiving: The Psychology of Visual Illusions. Lawrence Erlbaum, Hillsdale, NJ.
- Coren, S., Ward, L.M., Enns, J.T., 1999. Sensation and Perception. Hartcourt Brace College Publishers, Fort Worth, PA.
- Cott, H.B., 1940. Adaptive Colouration in Animals. Methuen, London.

- Endler, J.A., 1978. A predators view of animal colour patterns. *Evol. Biol.* 11, 319–364.
- Endler, J.A., 1987. Predation-light intensity and courtship behaviour in *Poecilia reticulata* (Pisces: Poeciliidae). *Anim. Behav.* 35, 1376–1385.
- Endler, J.A., 1992. Signals, signal condition, and the direction of evolution. *Am. Nat.* 139, 125–153.
- Endler, J.A., 1993. Some general comments on the evolution and design of animal communication systems. *Philos. Trans. Roy. Soc. B* 340, 215–225.
- Endler, J.A., Basolo, A.L., 1998. Sensory ecology, receiver biases and sexual selection. *Trends Ecol. Evol.* 13, 415–420.
- Enquist, M., Arak, A., 1994. Symmetry, beauty and evolution. *Nature* 372, 169–172.
- Enquist, M., Arak, A., 1998. Neural representation and the evolution of signal form. In: Dukas, R. (Ed.), *Cognitive Ecology. The Evolutionary Ecology of Information Processing and Decision Making*. The University of Chicago Press, Chicago, pp. 21–87.
- Enquist, M., Johnstone, R., 1997. Generalization and the evolution of symmetry preferences. *Proc. R. Soc. London Ser. B* 264, 1345–1348.
- Ghirlanda, S., Enquist, M., 1998. Artificial neural networks as models of stimulus control. *Anim. Behav.* 56, 1383–1389.
- Godfrey, D., Lythgoe, J.N., Rumball, D.A., 1987. Zebra stripes and tiger stripes—the spatial-frequency distribution of the pattern compared to that of the background is significant in display and crypsis. *Biol. J. Linn. Soc.* 32, 427–433.
- Guilford, T., 1990. The evolution of aposematism. In: Evans, D.L., Schmidt, J.O. (Eds.), *Insect Defences Adaptive Mechanisms and Strategies of Prey and Predators*. State University of New York Press, New York, pp. 23–61.
- Guilford, T., Dawkins, M.S., 1991. Receiver psychology and the evolution of animal signals. *Anim. Behav.* 42, 1–14.
- Haber, R.N., 1965. The effect of prior knowledge of stimulation on word recognition processes. *J. Exp. Psychol.* 69, 282–286.
- Haber, R.N., Hershenson, M., 1965. Effects of repeated brief exposures on growth of a percept. *J. Exp. Psychol.* 69, 40–46.
- Hanson, H.M., 1959. Effects of discrimination training on stimulus generalisation. *J. Exp. Psychol.* 58, 321–333.
- Hinde, R.A., 1970. *Animal Behaviour. A Synthesis of Ethology and Comparative Psychology*. McGraw-Hill Kogakusha, Ltd., Tokyo.
- Hurd, P.L., Wachtmeister, C.-A., Enquist, M., 1995. Darwin's principle of antithesis revisited: a role for perceptual biases in the evolution of intraspecific signals. *Proc. R. Soc. London Ser. B* 259, 201–205.
- Johnstone, R., 1994. Female preferences for symmetrical males as a by-product of selection for mate recognition. *Nature* 372, 172–175.
- Järvi, T., Sillén-Tullberg, B., Wiklund, C., 1981. The cost of being aposematic. An experimental study of predation on larvae of *Papilio machaon* by the great tit *Parus major*. *Oikos* 36, 267–272.
- Kamo, M., Kubo, T., Iwasa, Y., 1998. Neural network for female mate preference, trained by a genetic algorithm. *Philos. Trans. R. Soc. London Ser. B* 353, 399–406.
- Lieske, E., Myers, R., 1994. *Coral Reef Fishes. Indo-Pacific & Caribbean*. HarperCollins, London.
- Mackintosh, N.J., 1974. *The Psychology of Animal Learning*. Academic Press, London.
- Møller, A.P., Swaddle, J.P., 1997. *Asymmetry, Developmental Stability, and Evolution*. Oxford University Press, Oxford.
- Morton, E.S., 1975. Ecological sources of selection on avian sounds. *Am. Nat.* 109, 17–34.
- Mottram, J.C., 1916. An experimental determination of the factors which cause patterns to appear conspicuous in nature. *Proc. Zool. Soc.* 1916, 383–419.
- Ortolani, A., 1999. Spots, stripes, tail tips and dark eyes: predicting the function of carnivore colour patterns using the comparative method. *Biol. J. Linn. Soc.* 67, 433–476.
- Phelps, S.M., 2001. History lessons: a neural network approach to receiver bias and the evolution of communication. In: Ryan, M. (Ed.), *Annuran Communication*. Smithsonian Institution Press, London, pp. 167–182.
- Rothschild, M., 1964. An extension of Dr. Brower's theory on bird predation and food specificity, together with some observations on bird memory in relation to aposematic colour patterns. *Entomologist* 97, 73–78.
- Ruxton, G.D., 2002. The possible fitness benefits of striped coat coloration for zebra. *Mammal. Rev.* 32, 237–244.
- Ryan, M.J., 1998. Sexual selection, receiver biases, and the evolution of sex differences. *Science* 281, 1999–2003.
- Schuler, W., Hesse, E., 1985. On the function of warning coloration: a black and yellow pattern inhibits prey-attack by naive domestic chicks. *Behav. Ecol. Sociobiol.* 16, 249–255.
- Staddon, J.E.R., 1975. A note on the evolutionary significance of “supernormal” stimuli. *Am. Nat.* 109, 541–545.
- Stewart, I., Golubitsky, M., 1992. *Fearful Symmetry. Is God a Geometer?*. Penguin Books Ltd., London.
- ten Cate, C., Bateson, P., 1989. Sexual imprinting and a preference for “supernormal” partners in Japanese quail. *Anim. Behav.* 38, 356–358.
- Thayer, A.G., 1918. *Concealing Colouration in the Animal Kingdom*. Macmillan, New York.
- Tinbergen, N., 1951. *The Study of Instinct*. Clarendon Press, Oxford.
- Turing, A.M., 1952. The chemical basis of morphogenesis. *Philos. Trans. R. Soc. London B* 237, 37–72.
- Wells, F.E., Clayton, W.B., 1993. *Sea Slugs of Western Australia*. Western Australian Museum, Australia.
- Zahavi, A., 1975. Mate selection—A selection for a handicap. *J. Theor. Biol.* 53, 205–214.
- Zahavi, A., 1977. The cost of honesty (Further remarks on the handicap principle). *J. Theor. Biol.* 67, 603–605.