

Ritchie, G. and Kirby, S. (2006). Modelling the transition to learned communication: an initial investigation into the ecological conditions favouring cultural transmission. In Cangelosi, A., Smith, A., and Smith, K., editors, *The Evolution of Language: Proceedings of the 6th International Conference on the Evolution of Language*. World Scientific.

MODELLING THE TRANSITION TO LEARNED COMMUNICATION: AN INITIAL INVESTIGATION INTO THE ECOLOGICAL CONDITIONS FAVOURING CULTURAL TRANSMISSION

GRAHAM R. S. RITCHIE and SIMON KIRBY

*Language Evolution and Computation Research Unit,
University of Edinburgh, Edinburgh, UK
graham.ritchie@ed.ac.uk / simon@ling.ed.ac.uk*

Vocal learning is a key component of the human language faculty, and is a behaviour we share with only a few other species in nature. Perhaps the most studied example of this phenomenon is bird song which displays a number of striking parallels with human language, particularly in its development. In this paper we present a simple computational model of bird song development and then use this in a model of evolution to investigate some of the ecological conditions under which vocal behaviour can become more or less reliant on cultural transmission.

1. Introduction

One of the most unusual characteristics of language, when compared to many of the other communication systems found in nature, is the extent to which it relies on vocal signals transmitted culturally rather than genetically. This is of considerable interest as other modelling work has demonstrated the role that cultural transmission, via ‘iterated learning’, may play in explaining many prominent features of human languages, e.g. the emergence of compositional syntax (e.g. Brighton, 2002), regular and irregular word forms (Kirby, 2001), and dialects (Livingstone, 2002). The evolution of learning can therefore be seen as a key transition in the evolution of human language.

Vocal learning is a comparatively rare evolutionary development, it appears to have only evolved in three groups of mammals: humans, bats and cetaceans, and three groups of birds: songbirds, hummingbirds, and parrots (Jarvis, 2004). Of these, the development of bird song and human language have a number of striking similarities, e.g. both nestlings and human babies have a critical period for learning, both rely on auditory feedback for normal development, and both exhibit a form of early babbling (known as subsong in birds) (Doupe & Kuhl, 1999). This suggests that there may be strong epigenetic constraints on the evolution of a learned vocal system (Jarvis, 2004), and so studying the evolution of learning in bird song may help us to elucidate possible ecological factors which played a role in the transition to learned communication in our own species.

In this paper we use a computational model to investigate the possible role of two very simple ecological conditions which we think may affect the transition

to learning; namely the reliability of cultural transmission, and the stage of life at which communication is required.

2. The auditory template model of song development

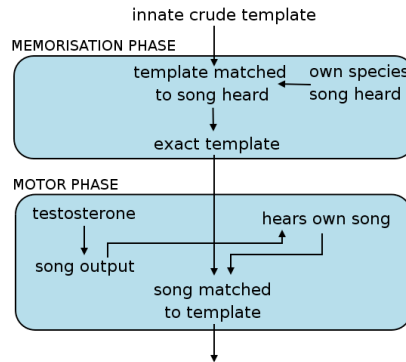


Figure 1. The auditory template model of song development. Figure after (Catchpole & Slater, 1995).

The song learning behaviour of many different species of the oscine passerines has been extensively studied, for an introduction see the reviews in (Catchpole & Slater, 1995) and (Marler & Slabbekoorn, 2004). The exact pattern of song development varies greatly among different species, but in attempting to capture the general features, bird song biologists have developed what is known as the ‘auditory template model’ of song learning, depicted in figure 1. This model posits two distinct phases to song learning; an early memorisation phase in which songs that are heard as an infant which are recognised as conspecific by an innate ‘crude template’ are memorised, and a later motor phase when song production is trained to produce songs that match the learned template. This behaviour can be contrasted with the sub-oscine passerines which appear to have a largely innately specified song, and will develop normal song production without hearing conspecific song and without auditory feedback.

3. A simple computational model

We take this model as our inspiration and develop a computational model of the two stages of learning in bird song, described in the preceding two sections. We then use this model to investigate some conditions under which song perception and production can come to be increasingly influenced by cultural transmission.

3.1. Phase 1: Observational learning

To model the memorisation phase of song learning we hypothesise a module which we term the Species Recognition Device (SRD). This is intended to model the auditory biases birds appear to show towards conspecific song.

We model the SRD as a note transition matrix which defines the transition probabilities between every available note (or song element)^a. We assume that the notes are fixed and identical for every agent in the simulation, and the number of notes used here was 6. We realise that this is unrealistic and that many species learn the form of song elements from their tutors as well as the element sequence. We also realise that element transitions or sequence are not the only cues birds use to identify conspecific song. However, a note transition matrix provides us with a simple and computationally tractable model of these sorts of biases.

Each agent in the model has ‘genes’ which code for an innate SRD, this is intended to model the ‘crude template’ as described above. An agent uses its SRD to categorise songs it hears as either conspecific or not by comparing the note transitions in the song with the transition probabilities in the matrix.

Such a matrix can be more or less biased to a particular song-type, if all the probabilities in the matrix are equal then the matrix has no preference to any particular song, while if each row has exactly one high probability transition, the matrix is maximally biased to one particular song. We can measure this bias by calculating the Shannon entropy for each transition distribution, and we can measure the preference of a matrix for a particular song by comparing the transitions found in the song and the probabilities in the matrix. We have used these measures of matrix preference and bias in earlier work (Ritchie & Kirby, 2005), and the reader is referred there for a more detailed definition.

An agent’s adult SRD is also subject to being altered by songs heard in early life, we model this by ‘exposing’ each agent to 100 songs from its environment and getting it to select the ones preferred by its innate SRD (crude template). The note transitions in the songs that are selected at this stage are then reinforced in the agent’s SRD to produce the agent’s adult SRD, or ‘exact template’. The degree to which an agent’s SRD is modifiable by songs heard in early life is determined by genes which code for the agent’s SRD plasticity (SRDP), this will be a value between 0 and 1, with 0 meaning the innate SRD is entirely fixed, and 1 meaning that the agent relies only on songs heard early in life to construct its adult SRD.

3.2. Phase 2: Reinforcement learning

The SRD as described in the previous section models an agent’s sensory biases (or lack thereof) to a particular song-type. We also require a model of song production. We also model this as a note transition matrix^b, but here the probabilities determine the probabilities of singing one note after another. We call this the Song Production Device (SPD). Just as for the SRD, an agent encodes innate biases for

^aWhile we implement the SRD as a note transition matrix here, we hope that this component could be modelled in many different ways, e.g. as a neural net with the initial weights specified genetically.

^bAgain, we hope that the SPD component could be modelled in a number of different ways, not necessarily using the same mechanism as for the SRD.

its SPD in its genes.

To model plasticity in the production mechanism, we allow the SPD to be trained by reinforcement learning using the agent's SRD as a critic, using a very simple learning algorithm. This is intended to model the process by which a bird uses its memorised exact template to guide its vocal development. As for the SRD, the degree to which the adult SPD is allowed to be influenced by learning, the SPD plasticity, (SPDP), is determined genetically. If the plasticity is 0 then the SPD is not influenced at all by the learning procedure described below, higher values mean the SPD becomes increasingly influenced by learning.

The SPD is trained by getting the agent to produce a song and then to 'listen' to this song with its adult SRD, if a note transition in the song is 'accepted', i.e. has a high probability in the SRD matrix, that transition's probability is increased slightly in the SPD. This process is repeated 250 times, after which the agent's SPD is said to have 'crystallised' and will not change again in the agent's lifetime^c.

3.3. Determining fitness

We define an agent's fitness as its ability to recognise and be recognised by conspecifics. This seems a reasonable model of one of the main pressures acting on song (Catchpole & Slater, 1995), although there are of course many other pressures acting on song in the wild (e.g. sexual selection for variation, adaptation to the local acoustics etc.), and we hope to model some of these in future work.

To calculate an agent's fitness we perform 250 fitness trials. In each trial we get the agent to produce a song using its crystallised SPD and we then randomly select another member of the population and check that this second agent correctly recognises the song using its adult SRD. We also get the second agent to produce a song and check that the first agent correctly recognises the song. Every correct recognition means that the agent's fitness is incremented by 1.

Defining fitness in this way means that there is a strong selection pressure for the agents to develop and maintain a stereotypical and easily recognised species-specific song. As the SRD is modelled as a note transition probability matrix, this corresponds to a matrix with a single high probability transition for each individual note. In short, in this environment it is adaptive to have strongly biased matrices.

3.4. Overall model design

The overall model works with an evolving population of 100 agents. As we want to investigate how a genetically specified song can come to be learned we initialise the agents' innate SPD and SRD genes to one particular song "abcd", and the plasticity genes to 0. This means that the population will start off receiving maximal

^cUnfortunately we do not have space to describe the learning algorithm in detail here. Further details are available upon request, and will be described more fully in future work.

fitness values and any mutations that degrade an agent's ability to sing and recognise conspecific song will be selected against. Each agent in each generation then goes through the following 'life stages':

Birth The agent's innate SRD and SPD, along with its SRDP and SPDP, are decoded from its genes.

Development Each agent is exposed to the songs of the previous generation, and picks those which will be used for learning using its innate SRD. The agent then goes through the two stages of learning described above to give them their adult SRD and crystallised SPD.

Adulthood The agents are tested in 250 fitness trials as described above to see how many times it can correctly recognise a bird of its own species and how many times its song is correctly recognised by a bird of its own species. These values are summed to give a bird's fitness score.

Reproduction Parents from the population are selected probabilistically according to their fitness score and their genes are recombined and subject to a low mutation rate to produce new child agents^d.

Death Each bird in the population is sampled 5 times and the resulting songs are stored for the next generation to learn from. All of the current birds in the population are removed and their children become the new population.

We repeat this process over many generations and record various measures over the course of a run.

4. Experiments

In this initial investigation we only model two very simple ecological conditions:

Environmental reliability For the first experiment we vary the reliability of the environment, that is the degree to which the previous generation's songs are faithfully recorded and then passed on to the new generation to learn from. We have two conditions: a reliable environment where we keep 80% of the previous generation's songs, and an unreliable environment where we keep only 20% of the previous generation's songs. The remaining songs are randomly generated songs which use the same notes and are constrained to within the same length as the agents' songs. This intended to model heterospecific song or other extraneous sounds in the birds environment.

Timing of song requirement In the first experiment we only test the bird's fitness after learning has taken place, in this experiment we also check the bird's fitness *before* learning. This intended to model a possible environment in which song is

^dThis is implemented with a standard genetic algorithm (GA), using tournament selection, a crossover rate of 0.7 and a mutation rate of 0.01. Mutation is modelled by simply replacing the gene that is to be mutated with a uniform random number between 0 and 1.

required immediately after birth as well as later in life.

5. Results

We provide results for each of the three different conditions described above in figure 2. The measures shown in each are the population average fitness, SPDP, SRDP, SPD change and SRD change. The SPD and SRD change are simply the absolute difference of the bias value of the innate and adult matrices (as discussed in section 3.1 above). We measure this as well as the plasticity values as the plasticity values can vary *without* a correlated variation in the change values (as demonstrated in figure 2c).

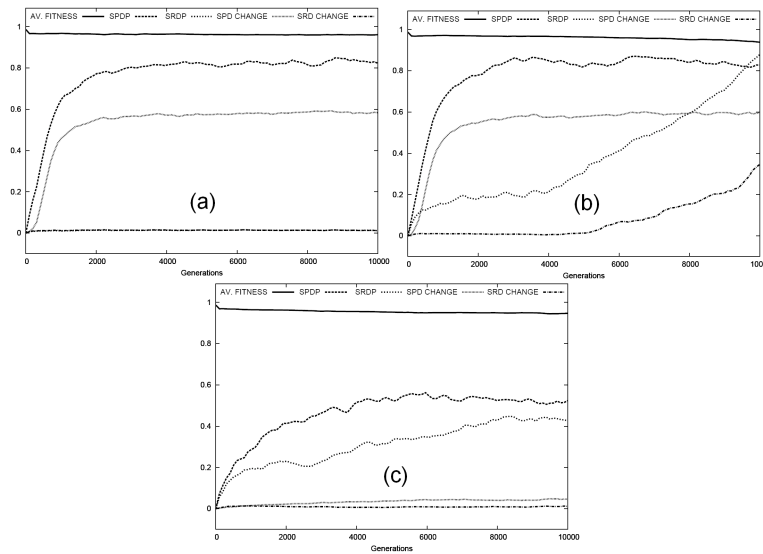


Figure 2. Results for the three different environments. The X-axis in each graph is the number of generations, set to 10000 for all results shown here. The Y-axis in each graph measures the population average fitness, SPDP, SRDP, SPD change and SRD change for each different condition. Graph (a) shows results for an unreliable environment where only 20% of the previous generation's song are faithfully passed on. Graph (b) depicts a reliable environment where 80% of the songs are passed on. Graph (c) shows results for a reliable environment in which the agents' fitness is checked both before and after learning. These results are the averages of 10 separate runs for each condition with a different random number generator seed for each. We have smoothed the graphs to allow us to better see the overall trends.

In all of the conditions we found that fitness stayed fairly fixed throughout all of the runs. However, the degree to which song remained being transmitted genetically depended on the environment, as demonstrated by the different values of SRDP and SRD change at the end of each simulation.

In the unreliable environment the population cannot count on hearing conspecific song as infants. The agents therefore have to keep transmitting their song

genetically, as demonstrated by the much lower SRD change and SRDP at the end of the run in figure 2a. In contrast, in the reliable environment shown in figure 2b, towards the end of the runs the population begins to transmit their song culturally as demonstrated by the coincident rise in the population's SRDP and SRD change.

In both experiments, however, the SPD change and SPDP quickly rise, indicating that the SPD is always being trained using the adult SRDs and the reliability of the environment appears to have no bearing on this. As long as the adults can construct a faithful copy of the species song in their SRDs as a result of either cultural or genetic transmission, it can always be used to train the SPD, and so there is no pressure for the copy of song stored in the SPD to be transmitted genetically and mutation pressure quickly erodes the genetic copy.

Figure 2c show results when the timing of song requirement is changed, where we test an agent's fitness both before and after learning. The SPD and SRD change values stay low throughout the run, demonstrating that SPD and SRD copies of song remain genetically transmitted throughout the run. The average SPDP and SRDP values drift to around 0.5 as there is no selection pressure acting to maintain these at any particular value.

6. Discussion

The results described here predict two simple environmental conditions which could affect the transition to a learned communication system; the reliability of the cultural environment, and stage of life at which communicative behaviour is required. These conditions seem fairly widely applicable and it seems reasonable that these conditions may have played a role in the transition to increased reliance on learning in human communication as well.

We think that this model also provides an interesting case study of the interaction of genetic and cultural transmission and phenotypic plasticity. We see that where the environment is reliable enough, and a learning mechanism is available to the population, the genes need not code for a song explicitly as an agent can rely on obtaining a copy of the 'correct' song via cultural transmission. Cultural transmission can thus, in some conditions, be seen as a *masking* force (Deacon, 2003) on genetic transmission, with a similar end result to that we found in earlier work (Ritchie & Kirby, 2005) for rather different environmental conditions.

Another interesting result is that in all of the experiments described here the agents come to rely solely on their auditory copy of song (in the SRD) to guide later production behaviour. We feel that this again represents a form of genetic parsimony, as it seems rather inefficient for an agent to store two 'copies' of their song genetically, even though these copies are likely to be represented in rather different ways; one being a sensory and the other a motor mechanism. Nevertheless, if there is enough phenotypic plasticity to allow these to interact, and if the genetic 'cost' of this plasticity is lower than the cost of encoding a song genetically, we see that even in the unreliable environment the agents rely on only

their auditory copy, but need it always be this way round? In the case of bird song it seems so, as a bird only needs to produce a song when it is sexually mature while it needs to be able recognise conspecific song earlier. This means that the song recognition system should be more genetically constrained than the song production system, which seems to match the biological data.

While this may be true of bird song it is not so clear for human language as children become capable talkers well before puberty. In future work we would like to relax some of the assumptions built into the current model with regard to the timing of each of the learning phases and allow this to be under genetic control. The two ecological conditions we discuss here are the simplest relevant condition we could think of, and we would also like to model other relevant ecological conditions, such as sexual selection pressure, to see what role these may play in conjunction with the conditions investigated here.

References

- Brighton, H. (2002). Compositional syntax from cultural transmission. *Artificial Life*, 8(1), 25–54.
- Catchpole, C. K., & Slater, P. J. B. (1995). *Bird song: Biological themes and variations*. Cambridge University Press.
- Deacon, T. (2003). Multilevel selection in a complex adaptive system: the problem of language origins. In B. Weber & D. Depew (Eds.), *Evolution and learning: the baldwin effect reconsidered* (pp. 81–106). MIT Press, Cambridge, MA.
- Doupe, A. J., & Kuhl, P. K. (1999). Birdsong and human speech: Common themes and mechanisms. *Annual Reviews of Neuroscience*, 22, 567–631.
- Jarvis, E. D. (2004). Brains and birdsong. In P. Marler & H. Slabbekoorn (Eds.), *Nature's music: The science of birdsong* (pp. 226–271). Academic Press Inc. (London) Ltd.
- Kirby, S. (2001). Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Journal of Evolutionary Computation*, 5(2), 102–110.
- Livingstone, D. (2002). The evolution of dialect diversity. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 99–118). London: Springer Verlag.
- Marler, P., & Slabbekoorn, H. (2004). *Nature's music: The science of birdsong*. Academic Press Inc. (London) Ltd.
- Ritchie, G., & Kirby, S. (2005). Selection, domestication, and the emergence of learned communication systems. In *Proceedings of AISB 2005: Social intelligence and interaction in animals, robots and agents*.