

Modelling Musical Behaviour in a Cultural-Evolutionary System

Oliver Bown and Geraint A. Wiggins

Centre for Centre for Cognition, Computation and Culture,

Department of Computing,

Goldsmiths College,

New Cross, London, SE14 6NW, UK

{o.bown, g.wiggins}@gold.ac.uk

Abstract

The computational study of musical behaviour, as with creativity, is traditionally couched in psychological terms, but in both cases a social perspective has also taken root. We present an approach to studying musical behaviour through the simulation of cultural evolutionary systems. We discuss the need for this approach and some of the issues involved in constructing such a simulation, and present observations from our first trials, which suggest ways in which evolving systems of musical normativity can be examined. These initial observations will form the basis for future models.

1 Introduction

According to Csikszentmihalyi, creativity needs to be seen in a social context: “what we call creative is a phenomenon that is constructed through an interaction between producer and audience. Creativity is not the product of single individuals, but of social systems making judgements about individuals products” [Csikszentmihalyi, 1999, p. 314]. This *systems theoretic* perspective brings evaluation within the creative process. We may say that there are objective aspects to value, but we cannot talk about value *out of context*.

The research presented here is based on a scientific interest in the social dimensions of creativity, and our focus is on the dynamic relationship between value and experience in a society of learning individuals rather than on the production of novelty. We aim to understand how value is constructed in an artificial system in which creative products cannot be said to have any objectively assigned value, and we assume that musical creativity fits into this category. Our long term aim is to see what kind of relations can exist between major entities in an evolving cultural system: genetic structure; cultural norms such as musical style; and social groups and relationships.

A number of emerging theories on the origins of music emphasise the role that musical behaviour plays in human social life and cognitive development [Brown *et al.*, 1999]. Richman [1999], for example, describes language and music as simultaneously emerging from a primitive state in which evolving ‘nonsense’ vocalisations played a role in social grooming between individuals. Our central point of investigation with respect to music is the notion that musical behaviour mediates a relationship between musical normativity and social organisation.

In the present paper we wish to indicate how such relationships may be studied by considering one simulation, a single example of a simulation framework which will be used for further studies. In this model multiple individuals interact by sharing music, moving in a metric space, and assigning status to each other, as well as dying, reproducing and adapting their behaviour through experience. Since this is work in progress, we are neither in a position to precisely determine testable hypotheses nor present statistically verified results: we are interested in generating hypotheses from the observations of initial exploratory simulations. Thus we present a single run of a single instance of the system, which was chosen for its interesting behaviour from a series of runs with various parameters. We treat this as an indication of the type of analysis that is possible in this area, and the problems and potential benefits we are likely to encounter in future work.

2 Methodology

Human social systems, resting as they do on intelligent human behaviour, are absurdly complex. With this in mind we expect to have to produce a model that is convoluted in its design and quantity of variables, but nevertheless fails to approach ‘realism’. We adopt the methodology of Di Paolo *et al.* [2000], treating simulations of complex systems as *opaque thought experiments*. That is, they are “valuable tools for re-organising and probing the internal consistency of a theoretical position” [Di Paolo *et al.*, 2000, p. 497] in cases where a simulation is required to help think through the complexity (hence opacity) of the system in question.

Di Paolo *et al.* [2000] propose a three stage approach to complex simulation models which we quote in full:

1. *Exploratory phase*: After the initial simulation model is built, explore different cases of interest, define relevant observables, record patterns, re-define observables or alter model if necessary.
2. *Experimental phase*: Formulate hypotheses that organise observations, undertake crucial “experiments” to test these hypotheses, explain what goes on in the simulation in these terms.
3. *Explanatory phase*: Relate the organisation of observations to the *theories* about natural phenomena and the hypotheses that motivated the construction of the model in the first place, make explicit the theoretical consequences.

The present work is situated in the exploratory phase; we report on the basic patterns of behaviour observed in the first

version of the model and propose developments based on these observations.

3 Previous Work

Although there is a great deal of work in simulating the evolution of language, and in using evolution in creative musical activities, there has been far less work in the evolution of musical behaviour in nature. Werner and Todd [1997] consider the conditions for musical diversity to emerge from coevolutionary processes in a model of sexual selection in birdsong. Male birds' songs are represented as strings of 32 notes taken from a selection of 24 possible note values. Female birds judge the males' songs using a table of probabilities for transitions from one note to the next. Male songs and female transition tables are encoded together into a combined genotype which is recombined and mutated during reproduction. At each generation, each female listens to n males chosen at random and mates with their favourite male to produce one child; thus each female is the mother of exactly one child, but males may father a number children (or none at all) depending on their success. Werner and Todd vary three principal elements: the number n of males that each female samples before choosing a mate; whether or not the system is coevolutionary; and the method that the female uses to evaluate the male song. For the latter the main distinction is between favouring songs that satisfy the expectations of the transition table, and those that contradict the expectations; *i.e.*, that have some element of surprise. Werner and Todd examine the effects of these parameters on the synchronic and diachronic diversity of the male songs throughout the simulation. They find that the use of coevolution, with a small sample size, leads to sustained song diversity over time and across the population.

Our focus of attention differs from Werner and Todd's in two respects. First of all, due to our interest in social organisation we require a *space* in which individuals may be organised, which may favour some interactions over others. Di Paolo [1997] has shown that spatial versions of game-theoretic models have new properties that result from the effect of spatial organisation on environment and on behaviour.

Second, we focus on the role of a social environment in looking at the relationship between musical behaviour and genetic selection. Whilst sexual selection is a serious and credible theory for the origins of music, we do not think it offers a complete account of music in human evolution, which is rich in non-sexual social behaviour. In the present model we do not consider genetic evolution at all, with the aim of understanding how genetically identical individuals take on *culturally* determined behavioural traits, which can be achieved either through some form of adaptation, or through a direct connection between environment and behaviour. We note that these two factors have the potential to provide synchronous diversity for free: in cases where individuals randomly vary but never interact they are likely to diverge in behaviour. This leaves little scope for comparison with Werner and Todd's model.

4 Towards a Simulation of Musical Behaviour in a Cultural Evolutionary System

In this section we present the first instantiation of our simulation framework, beginning with the simplifications and as-

sumptions made. We define simplifications as those conditions that we know not to be true, but that we expect to be able to ignore, and assumptions as those conditions we take to be true for the purpose of a particular model. The assumptions made here are admittedly speculative, based on informal observations of human musical behaviour rather than on empirical research. It is likely that in future work the simplifications listed below will largely remain fixed while the assumptions will be varied. Due to the exploratory nature of this work, we do not detail any expectations for this model, except to reiterate that we are interested in the possibility of feedback processes between social organisation and musical behaviour.

4.1 Simplifications

Our model's greatest simplification is the most methodologically tricky: although we aim to look at musical behaviour as an interactive participant in the complex dynamics of social life, we are obviously unable to model 'the rest of social life'. Our first simplification, then, is that *musical behaviour can be modelled independently from the 'rest of social life'*. Other simplifications are more clear-cut: we limit the scope of musical behaviour by the definition that *musical acts are single events performed by one individual and perceived by one or more others*. This rules out ensemble performance, and effectively de-temporalises musical experience in order to focus on the macroscopic evolution of musical systems. However, it is our intention to maintain the temporality of music perception in the cognitive aspects of the model. Next, following Werner and Todd, *the musical representation is highly simplified* for the sake of transparency. Indeed we copy Werner and Todd's Markov representation in this instance. Next, we model the space of social connections as a two-dimensional Euclidian space. That is, *we reduce social proximity to physical proximity*. From our first simplification, we have necessarily detached musical interaction from ingrained social ties such as kinship, leaving only idealised individuals with complete social mobility. We tie down this unlimited mobility by situating individuals on a 2-D plane, representing the fact that 'friendship' is often transitive. Finally, *we treat status as a universally visible attribute of individuals which results directly from the positive or negative effects induced in others through communication of signals*.

4.2 Assumptions

Our central assumptions for the initial model are as follows:

- musical preference is determined passively through experience;
- attention to a source of music is determined by the social proximity to and status of that source;
- following a musical experience, individuals adjust their position in a social space according to some combination of their evaluation of that experience, the status of the source of the experience, and their own status;
- individual lifespans and reproductive success are independent of musical behaviour, social context and status;
- individual preferences are based on positive expectation, but negative responses can also be induced by over-familiarity.

4.3 Simulation Settings

We have already described the musical model appropriated from the work of Werner and Todd. We have adapted this model to incorporate learning in the simplest way possible; for each transition that the individual is subjected to we increment the appropriate cell in the individual's transition table by a small amount (fixed at 0.1 for these simulations), and then re-scale the probabilities in the relevant row to sum to 1. All individuals begin with a flat transition table; *i.e.*, each transition is equally likely. Individuals generate 'songs' using the probabilities from the transition table and a uniform probability distribution.

Our evaluation method is inspired by Saunders [2001], who models the behaviour of 'curious' design agents using the effect of the Wundt Curve, which maps subjective familiarity to value. For Saunders' design agents, the mapping assigns zero value to completely familiar percepts, increasing to a peak as familiarity decreases and then falling asymptotically to a negative value as familiarity decreases further. The curve embodies the intuitive principle that over-familiarity results in boredom, whilst unfamiliarity results in dislike, leaving a favourable zone somewhere in-between.

We employ a similar principle for our individuals' perceptual modules, using Werner and Todd's [1997] local-preference rule – which simply sums the expectation of each transition perceived – to return a positive score, but cutting off the score above a certain point (3.0) and returning a flat value of -1.5. Thus in our case value increases with familiarity up to a hard boundary, above which percepts are disliked. This choice will make more sense in the context of equation (2) below, which relates the evaluation to individual movements.

Individuals are updated synchronously in two stages: first each individual determines its *focus*: the individual to which it is going to listen. Each individual selects its focus from its nearest 5 neighbours, choosing 3 at random and evaluating their suitability according to a weighted combination of status and distance, as given by R_{xy} in (1). The individual with the highest score in this contest becomes the focus. Secondly, each individual evaluates the output of their focus, and then moves towards the focus by a distance determined by a weighted combination of its evaluation, its focus' status, and its own status (negative values are possible here) as given by M_{xy} in (2). This weighted combination is modified by a tanh function and scaled according to the current distance between the two individuals. Each evaluation is then added to the status of the focus.

$$R_{xy} = W_1^R D_{xy} + W_2^R S_y \quad (1)$$

$$M_{xy} = \alpha V_{xy} \tanh(W_1^M E_{xy} + W_2^M S_x + W_3^M S_y) \quad (2)$$

In these equations W_j^i are weights, D_{xy} is the distance between x and y , and S_x is the status of individual x , E_{xy} is the evaluation of individual y 's output by individual x , and V_{xy} is the vector from the position of individual x to that of individual y .

The values of W_1^R and W_2^R were fixed at -0.2 and 0.1 respectively. W_1^M , W_2^M and W_3^M were fixed at 0.2, 0.005 and -0.001 respectively. α was fixed at 0.1. These values were chosen based on a visual response to the behaviour of the system in earlier trials.

Individual lifespans were drawn from a normal distribution with mean 60 and standard deviation 10. When an individual

died it was removed from the population and a new individual was immediately put in its place with a slight displacement along each axis drawn from a normal distribution with standard deviation 1.0.

The simulation was run for 500 time steps. At the start of the simulation, 20 individuals were positioned within the space with a 2-dimensional normal distribution around the origin with standard deviation 10.0. Their integer ages were selected from a uniform distribution between 0 and 60. All beginning individuals' transition tables were initialised with equal probabilities throughout.

5 Exploration

In preliminary trials, clusters of individuals were seen to form rapidly with sizes close to, but never lower than the neighbourhood size. Positions of clusters were relatively stable, but close clusters would occasionally merge, die out, or drift. The evaluation rule and relative coefficients of (1) and (2) also strongly influenced the spatial clustering behaviour. We do not report on those effects here but instead focus on the relation between processes in a single representative case.

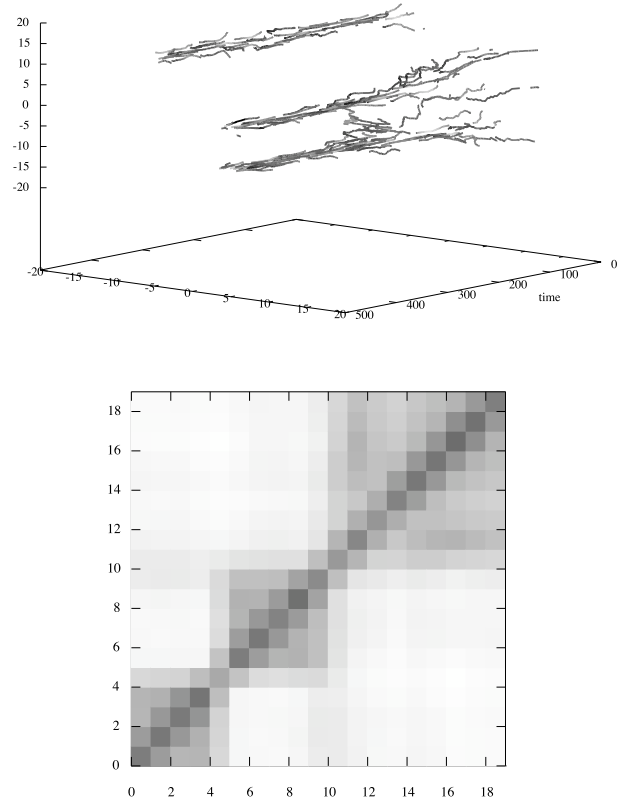


Figure 1: Simulation over 500 time steps showing: (top) spatial positions of individuals over time, with individual status values given by the line shading; (bottom) dissimilarity for all individuals' transition tables in the final population, with darker shades indicating greater similarity. Individuals were ordered according to proximity to a single individual in the top-left cluster.

Figure 1 (top) shows the trajectories of each individual

in the population over time for the representative case, with grayscale shading indicating the status of individuals during their lifetimes. The lower figure shows a transition table similarity plot for the whole population at the end of the run, in which darker shading corresponds to closer similarity. The population has been ordered according to proximity to a single individual, and so the plot indicates the relationship between physical and behavioural clustering. Three distinct blocks of greater similarity can be seen, corresponding to the three spatial clusters in the top of Figure 1. This is to be expected since there is no communication of signals between clusters: as long as each cluster takes an unstable trajectory through transition table space, the clusters will diversify.




Figure 2: Change of average transition table for cluster 1 (bottom right of first graph in figure 1) over time.

Figure 2 illustrates that clusters do indeed move erratically through transition table space. However, horizontal streaks of dark and light shading show that the cluster settles into phases for short periods, apparently corresponding to the length of individual lifespans. Certain streaks are also clearly longer than individual lifespans. The larger of the three clusters was slightly more stable than the other two, suggesting that the necessarily small population sizes in this run may have served to limit the stability of clusters.

Individual lifespans followed a common pattern (top of Figure 3). Status increases to a peak and then decreases to a negative value. However, the maximum values of individual status vary greatly. The status curve appears to reflect the effect of the hard cut-off point on the preference function, suggesting that as high status individuals attract more attention and more individuals adopt the style of these influential individuals, the evaluation function of the listening individuals hits the cut-off point and they begin to award negative status. We expected this to be accompanied by dispersal of the cluster, but we assume that this failed to happen because there was a more dominant force in the movement calculation. The bottom of Figure 3 shows the make up of an individual's transition table over time. This suggests a possible pattern of change related to the general trend in status values.

In order to understand how stable cluster styles were, we looked at what would happen if one cluster was 'invaded' by an individual with a different transition table structure. Using exactly the same run we intervened at $T = 210$ to move a high-status individual from the larger cluster and relocate it to the centre of one of the smaller clusters. The move had

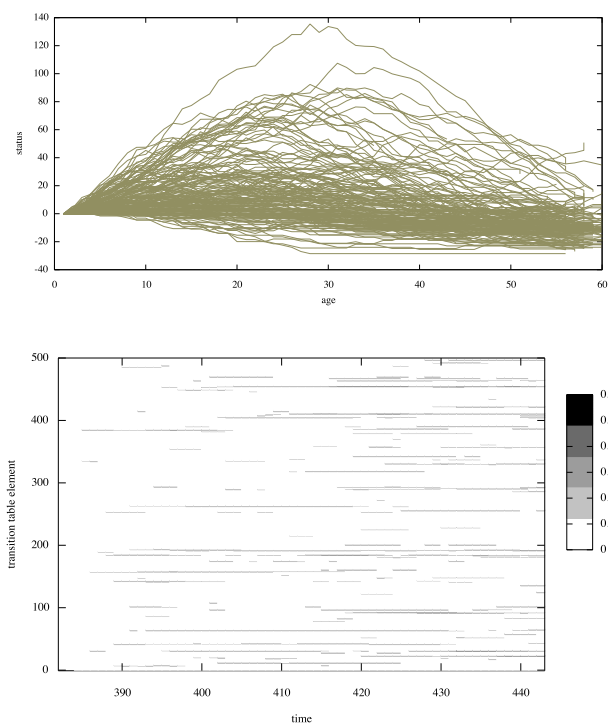


Figure 3: (top) Status against age for all individuals in the simulation, showing a clear trend for status to peak and then begin to fall. (bottom) example of change of transition tables in individuals over time.

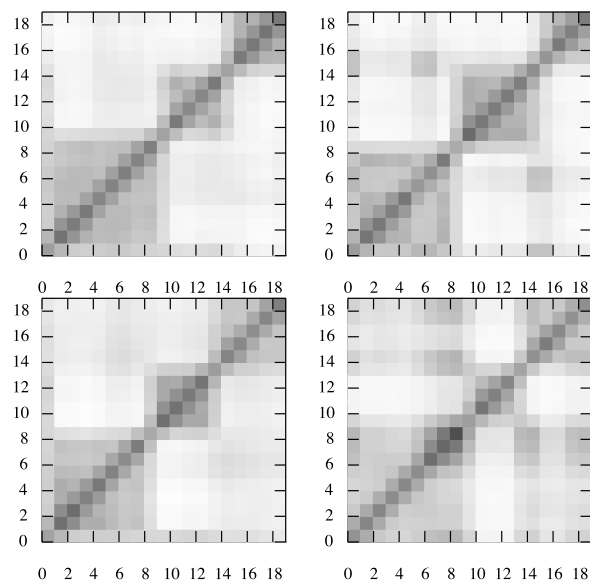


Figure 4: Similarity plots for the whole population at $T=201$ (top-left), $T=211$ (top-right), $T=221$ (bottom-left), and $T=236$ (bottom-right). An individual was moved at $T=210$ from the bottom-left to top-right cluster. After settling back into separate blocks, the individuals in the target cluster have grown more similar to those in the source cluster.

no apparent effect on the position of the cluster. Figure 4 shows four similarity plots constructed in the same way as the lower half of figure 1. The first plot, taken just before the move, at $T = 201$, shows a similar effect to the plot in figure 1. In the second plot, just after the move, at $T = 211$, the clear divisions between clusters have been slightly disrupted due to the presence of the foreign transition table in the target cluster. At $T = 221$ the divisions become clear again, but the similarity between the two clusters involved in the move has increased. After $T = 236$ this similarity increases further. We speculate that this is due to the influence of the invader.

6 Discussion and Future Work

The model as it stands constitutes an initial demonstration of how a computational approach to musical behaviour in the context of a convoluted coevolutionary system can prove informative. In this case spatial clustering is predominantly dependant on the neighbourhood size and not on feedback from other aspects of the system, as we initially set out to achieve. However, these initial observations suggest forms of analysis that can be used to aid a proper search of the parameter space and additional experimentation, as well as a number of directions in which the model could consequently be developed. We intend to move on to introduce genetic evolution into this system, but only after more detailed exploration of change on the sub-genetic time-scale.

Two of our most questionable simplifications were the musical model and the notion of a social space. With respect to the first, the present work suggests to us that it is necessary to abstract this further; defining musical traits, and predefining their behaviour, rather than relying on the behaviour of a model that is simple but fails to allow us to talk specifically about traits. At this stage we need a model which is more capable of explicitly indicating different behaviours. There is a wealth of candidate behaviour that we could ultimately investigate using specific musical models, such as proposed measures of objective aesthetics (see, for example, [Manaris *et al.*, 2003]).

With respect to the spatial simplification, we are also tempted to suggest that it is necessary to move towards greater abstraction for increased control. However, there is no obvious manner in which any other abstraction would be more beneficial, without losing the possibility for self-organisation. Since we are interested in the grouping of individuals, and content to assume that social relationships are transitive, the Euclidian spatial model appears to suffice. In this case the implementation of behavioural functions with respect to the social space is also key. In our model, movement in this space was implemented roughly in a physical sense. We will look into movement behaviour that is more flexible in the 2-D space that we are presently using, combining, for example, discrete leaps with continuous movement.

Turning to analysis, we have seen that the model exhibits a number of quantifiable features which we can ultimately consider with respect to real-world social behaviour: cluster stability in both physical and parametric (behavioural) space; synchronous and diachronic diversity; and the development of roles in individuals, *i.e.*, distribution of status within clusters, and correlations between this distribution and local patterns of change. Whilst clustering analysis is generally a straightforward process to automate, we have not so far been able to implement the tracing of cluster trajectories through

time autonomously, and in a way which correctly represents the division and combination of clusters. This is a necessary step in the automated search for cluster-related features. In the long term, a fixed repertoire of features will allow us to categorise model behaviour and perform searches across parameters for specific behaviours.

In terms of the social nature of creativity, we are keen to pursue interesting, dynamically changing evolutionary behaviour. We do not unquestionably assume that a dynamically changing social system indicates or drives creativity either in individual behaviour or in biological evolution, but our work is inherently concerned with this possibility.

Acknowledgments

The first author is supported by a Departmental Studentship from the Department of Computing at Goldsmiths' College, University of London. We would like to thank Peter M. Todd, Tim Blackwell and two anonymous reviewers for their helpful comments.

References

- [Brown *et al.*, 1999] Steven Brown, Björn Merker, and Nils L. Wallin. *The Origins of Music*. MIT Press, 1999.
- [Csikszentmihalyi, 1999] Mihaly Csikszentmihalyi. Implications of a systems perspective for the study of creativity. In Robert J. Sternberg, editor, *Handbook of Creativity*. Cambridge University Press, 1999.
- [Di Paolo *et al.*, 2000] E. Di Paolo, J Noble, and S. Bullock. Simulation models as opaque thought experiments. In *Artificial Life VII: The Seventh International Conference on the Simulation and Synthesis of Living Systems*, Reed College, Portland, Oregon, USA, 2000.
- [Di Paolo, 1997] E. Di Paolo. Social coordination and spatial organization: Steps towards the evolution of communication. In Phil Husbands and Inman Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life, ECAL97*. MIT Press/Bradford Books, 1997.
- [Manaris *et al.*, 2003] Bill Manaris, Dallas Vaughan, Christopher Wagner, Juan Romero, and Robert B. Davis. Evolutionary music and the zipf-mandelbrot law: Developing fitness functions for pleasant music. In *Lecture Notes in Computer Science, Applications of Evolutionary Computing*, LNCS 2611. Springer-Verlag, 2003.
- [Richman, 1999] Bruce Richman. How music fixed "non-sense" into significant formulas: On rhythm, repetition, and meaning. In Steven Brown, Björn Merker, and Nils L. Wallin, editors, *The Origins of Music*. MIT Press, 1999.
- [Saunders, 2001] R. Saunders. *Curious Design Agents and Artificial Creativity*. PhD thesis, Faculty of Architecture, The University of Sydney, 2001.
- [Werner and Todd, 1997] G. Werner and P. M. Todd. Too many love songs: sexual selection and the evolution of communication. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 434–443. Cambridge, MA: MIT Press/Bradford Books, 1997.