Automated generation of musical harmony: what's missing?

Geraint A. Wiggins
Music Informatics Research Group
Division of Informatics, University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, Scotland
geraint@ed.ac.uk

Abstract

Over the past thirty years, the automated generation of musical harmony has attracted much interest from Artificial Intelligence researchers. In particular, this interest has focussed on the problem of performing the task of *four-part vocal harmonisation.* The approach has generally been to capture the well-formedness rules for such a musical structure and to generate notes to conform to them.

In this paper, I argue that this approach is missing a fundamental point of music: creating music which listeners want to hear is not just a matter of filling in the notes according to the rules of harmony, but (much harder) of creating a musical structure which sets up and fulfils an *intention.* It is the intention which gives musical satisfaction to the listener, not the harmony alone.

I give examples of two systems which manipulate different aspects of musical intention, and suggest that, if Artificial Intelligence research is to offer anything to our understanding of human musical behaviour, this is the direction which future research must take.

1 Introduction

In this paper, I raise a question which seems to have been ignored by the existing Artificial Intelligence (AI) research into human musical behaviour, both in terms of listening and composing music: that of musical intention.

I place this in a context of a wide range of excellent AI work aiming to perform various compositional tasks, but none of which addresses the issue of large-scale musical structure *and the intention behind it*, in the mind of the composer and the performer.

The paper consists of an introduction to the concept of *musical structure*, (§2), a brief survey of some high-points of existing research in AI music systems (§3), followed by a more detailed outline (§4) of two AI music systems which address small parts of the larger problem addressed. I discuss the relevance of these case studies (§5) and conclude with a call for more research in this direction (§6).

Inevitably, in a discourse such as this, I will use more musical terms than I have space to define. The reader for whom this is a problem may find solace in publications such as that of Taylor (1996).

2 The notion of "Musical Structure"

2.1 What is musical structure?

Structure is what distinguishes music from noise. It exists in music on different *levels*, and it resides in various *places*. For example, we can think of a *sonata* (a particularly well-defined classical form, which specifies the relationships between several high level, large scale sections of a piece) as a musical structure; equally, at a much more local level, we can very usefully think of the progression from one chord to the next as structure.

Structure resides in various places because there are several aspects to communication of and by music, each of which may be ambiguous. So a composer may perceive a particular structure in her work, while the performer perceives another, possibly contradictory one; and the listener may have a different perception again. These apparent contradictions are often part of the joy of experiencing music, as there are always new ways of hearing to be discovered.

The human tendency to apply structure to heard sound can even mislead a listener into imposing a structure where there is none. If, for example, we hear a steady stream of notes, even if they are electronically produced to be mutually identical, we perceive a slight rhythm, either of two or three beat duration; this is imposed by our auditory processing system. On a different axis, if the notes of a melody are widely spaced in pitch, alternating between high and low registers, we perceive two separate melodic lines. This latter effect is called *auditory streaming* and is used to great effect in J. S. Bach's solo 'cello suites, allowing the single instrument to accompany itself.

2.2 Musical structure and time

The most important aspect of musical structure for my purposes here is the temporal relation between the different sections of a piece, at all levels. Music is fundamentally a time-based art, and it is the development (in some appropriate sense) of musical ideas, embodied as structures, which make a piece of music worth listening to.

This last point brings us to the difficult area of "musical meaning", carefully encapsulated in scare-quotes, as it is not immediately clear what "meaning" means in this context (Wiggins 1998a). I will sidestep this difficult question here, and would refer the reader to, for example, Lévi-Strauss (1964) and Meyer (1956) for a proper discussion. What is important is that there is *something*, which can cause affective and intellectual responses, directly measurable in terms of physiology, and reported consciously by listeners (Krumhansl 1997).

There are various aspects to whatever it is that is communicated by the music, and we can to a certain extent break down the signal into components. Perhaps best understood is the notion of *musical tension* (not to be confused with *dramatic tension*, which can also be communicated by music), which relates to the perceived level of harmonic completion of the musical phrase, to the local level of dissonance (in certain style-dependent ways), and to the rhythmic structure of the preceding passage.

One of the things reported by listeners as leading to a satisfying listening experience is an "ebb and flow" of (effects related to) musical tension. However, the notion of musical tension is not addressed explicitly in any of the AI composition systems surveyed in §3 below, and seems to be partially implicit in only one of them.

In answer to the question of my title, then: what is missing is an account of the underlying *intention* of music – the hidden structure that defines the ebb and flow of musical tension. Indeed, in most cases, what is missing is even an acknowledgement that such a thing exists.

Music without intention is not unlike language without semantics: it may be pretty and even entertaining. But it is not art.

3 Existing work on harmony generation

There is a substantial amount of existing AI work on harmony generation – enough, indeed, that a comprehensive survey would be impossible in a paper of this size. Therefore, in order to give a flavour of the work current in the field, I have selected representative examples of several different approaches, the aim being to demonstrate that there is little or no attempt in most to go beyond the correctness rules which characterise a particular style of harmony. For a fuller survey, the reader is referred to Papadopoulos and Wiggins (1999).

3.1 Grammar-based Systems

Grammar-based systems for generating music are founded on the assumption that the structure of music is more or less analogous with the structure of language, and therefore that it may be described by more or less the same mechanisms. This analogy can only be taken so far (and is misleading if over-stretched), but it is demonstrably useful in building models of harmonic style and convention.

A particular feature of grammatical music generators is that they usually aim to capture a particular style – in the case of the following examples, Blues, 17th-Century dance forms, and Indian traditional music, respectively.

Steedman (1996) uses a grammatical formalism (initially in the Chomskian style, and then, more convincingly, in Steedman's preferred Categorial style (Steedman 1999)) to capture the basic 12-bar blues style, with which most people are familiar from jazz and pop music. A particularly satisfactory aspect of the grammar is that it is able to capture very naturally the idea of *variations* – that is, that, given the basic 12-bar structure, one may alter parts of it in particular ways, while still remaining true to the original harmonic movement. Less satisfactory is that Steedman is required to impose meta-rules on his grammar which constrain the structures it produces to be of particular sizes – of course, linguistic grammars are not normally required to do this. While such extra constraint clearly is necessary to capture the Blues style, it is unsatisfying that the grammar formalism used cannot fully capture the nature of the musical structures, especially when such formalisms do exist – see below. Steedman refers to a "musical semantics", which is based on a notion of harmonic expectation founded in Longuet-Higgins' (1962) Tonal Harmony Space , and expounded more fully (though not completely) by Steedman (1994). This may, then, be something of a flavour of musical intention, though it is hard to see exactly how in the absence of more detail.

Ponsford, Wiggins, and Mellish (FC) use a statistical grammar-learning mechanism, based on 3- and 4-grams to capture the harmonic structure of a 17th-Century dance form, the Sarabande. The main grammar is derived from some 90 examples of sarabandes by various composers, and different sub-grammars are also given, to describe the styles of individual composers. Like Steedman's Blues grammars, this work suffers from a disparity between string-length requirements in language and music – the generated strings must be forced to be of the required length, a task which is onerous, since they are generated probabilistically, and so the process may be very time-consuming. Nonetheless, though, like Steedman's work, the grammars do capture the musical style rather well, they include no explicit notion of musical intention.

Bel and Kippen (1992) describe a system based on Pattern Grammars (Angluin 1980), which captures the *qa'ida* form of *tabla* drumming, which is traditional to India. Bel and Kippen take as their starting point that music is more akin to poetry than to ordinary language, and use the extra descriptive power of the pattern grammars to avoid the unsatisfactory meta-rules mentioned above. Interestingly, the output of Bel and Kippen's system was rated very highly by expert tabla drummers, even though it contains no explicit notion of "semantics" or intention at all.

Grammatical approaches to musical structure are summarised more fully by Sundberg and Lindblom (1991).

3.2 Knowledge Based Systems

There is in fact a fairly fine distinction between knowledge based systems and grammatical systems for music, since both approaches focus on the form being produced, and both are stated in terms of rules as to what can be juxtaposed with what. Here, I have simply categorised according to the perceptions of the original authors.

Perhaps the best known of all musical composition systems is the CHORAL programme (Ebcioglu 1988; Ebcioglu 1990), a system of some 350 rules, implemented in BSL, a logic programming language with "intelligent backtracking". CHORAL's expertise is in harmonising chorale (hymn) melodies in the style of J. S. Bach. Like most of the other systems covered here, CHORAL's rules make little mention of global musical structure or intention, though here at least, the nature of the melody is likely to impart significant clues to the KBS. In fact, it is possible that there is some encoding of intention in the search method used by CHORAL, but Ebcioḡlu's description of this is opaque, and it is hard to see what actually happens, though the harmonic rules themselves are stated clearly.

3.3 Genetic Algorithms and Genetic Programming

There has recently been a surge of interest in the application of evolutionary technology to artistic endeavour. Cynically, one might suggest that this is because genetic approaches to programming are particularly good at hiding the very question of the nature of musical creativity on which I am focussed here: we can begin with "canned" phrases from human musicians, and/or we can rely on human fitness assessors, and so bury the problem of creativity back where it came from.

The best known GA music system is GenJam (Biles 1994), a generator of jazz-style solos. Biles reports a successful range of output for GenJam, but this is perhaps unsurprising, since the program is an *interactive* GA – that is, its fitness function is the preference of a human listener. Aside from all the problems this brings (*e.g.,* very many solos to listen to; the inconsistency of human listeners), it is not very clear what scientific value the work has. In particular, it does not tell us *what makes the solos good* – it merely tells us which solos are (in Biles' opinion) good[1].

In an attempt to consider more formally how music might be generated by GA, Papadopoulos and Wiggins (1998), in jazz improvisation, and Phon-Amnuaisuk, Tuson, and Wiggins (1999), in chorale harmony, have both used *objective* fitness functions – that is to say, formally stated, computable functions, based on precepts drawn from the music theory and music

[1]The question of what constitutes "good" music is not one for a paper of this length – the reader is invited to choose his or her own favourite interpretation.

psychology literature. While both systems have been judged successful (the latter scraping a pass at 1st-year undergraduate level), they have both been criticised for lack of intention in their output – while both produce output which is technically and theoretically correct, it is not inspired music. The chorale system, in particular, is based on purely local musical rules, with no notion of larger scale structure or of musical intention.

As a result of the use of formal fitness functions, Phon-Amnuaisuk, Tuson, and Wiggins (1999) are able to study the nature of the harmonisation search space, and conclude that standard GAs are not well-suited to this kind of problem, due to the highly convoluted fitness profile produced. The point is that it is almost never possible to make local changes in a given harmonisation without significant knock-on effects, which means that search methods based on randomness are unlikely to be successful (Wiggins 1998b).

3.4 Constraint Satisfaction Systems

A more promising approach to the problem of generating harmony is based in the view of a search space of possibilities, limited by constraints. The chorale harmonisation problem (two part, as by Ovans (1992), or more complex, as by Pachet and Roy (1998)) may be viewed as a standard Constraint Satisfaction Problem, and treated as such. We can indeed be optimistic, since this view banishes the difficulties of the GA while still allowing us some prospect of coping with the enormous search space involved. A particularly promising version of the idea is that which abstracts away from individual notes, and considers chords (*i.e.,* localised lumps of harmony) as the initial level of reasoning (Pachet and Roy 1998). This is a step in the direction we need to begin working with musical intention, since this is usually thought of, at least initially, at the chordal level.

However, we are still left with the question of how to resolve ambiguity. Specifically, these constraint systems bear very large numbers of solutions, all of which are correct, in the sense of not breaking harmonic rules. However, many, if not most, of them would be rejected by expert musicians on grounds of "unmusicality", as was the output of Phon-Amnuaisuk's GA. What is not addressed here, then, is musical intention – that extra something which picks out a musically satisfactory harmonisation from an unsatisfying one.

3.5 Neural Networks

Finally, there have been attempts to capture harmonic progression via neural networks, such as that of Gang, Lehmann, and Wagner (1997). It has been suggested that this is an unnatural step to take, since

> "it is not entirely clear that net-based parallel techniques, which integrate over wide stretches of music with rather indefinite boundaries, are really appropriate to the task of interpreting chord sequences. This is particularly likely to be the case in assigning tonalities or chordal accompaniments to melodies, in which the transitions between tonalities which contribute to the identification of key seem to be quite abrupt and all-or-none in character."
>
> (Steedman 1994)

At least part of this criticism follows from the use of feedback in the network. This means that it is very difficult to analyse the information flow, and to understand and control the spread or locality of the information used. However, feedback is not the only way to model memory, and so some of Steedman's criticism can be avoided, as I shall show in §4.2, below.

In spite of the criticism, some success has been achieved, notably in the work of Barucha (1987). Again, however, the question of musical intention has been sidelined.

4 Two case studies in musical intention

In this section, I outline two experimental systems for music generation, one firmly symbolic and rule-based, the other connectist in character. Both of these systems are unusual in that they attempt to capture aspects of musical intention: in the former case, in creating dramatic tension in an educational environment, and in the latter, in directly modelling and reproducing harmony with specific tension levels.

4.1 HERMAN: a system for automated musical accompaniment

Background

Herman, named in honour (if inaccurately) of the well-known film composer, Bernard Herrmann, is an accompaniment system for an educational tool called GhostWriter. GhostWriter encourages young children to improvise ghost stories in an interactive virtual environment, which records their utterances for later editing and reconstruction into written stories. The environment is guided by a teacher, who is able to direct the construction of the story by means of various controls, and by communicating with the children.

One important aspect of GhostWriter is the achievement of *presence* for the participant children. Presence is that property of an artificially mediated experience whereby the subject becomes (temporarily) unaware of the fact of mediation – for example, presence is achieved when we are drawn into a cinema film or TV programme to the extent that we forget that we are not directly involved. More about GhostWriter and about presence are given by Robertson, de Quincey, Stapleford, and Wiggins (1998). One of the features of GhostWriter which contributes to presence is the background music, supplied to suit the nature of the situation currently being played out, by Herman.

HERMAN

Herman is a stochastic rule-based system which uses a database of musical knowledge, combined with "sensory" input (supplied by the teacher) about the situation being played out in GhostWriter. To do this, Herman combines a theory of harmony (Schoenberg 1978) and a simple theory of rhythm developed from poetic feet and from observation of the work of film composers, with a multi-dimensional control system, whose metric space is mapped onto a single dimension for simple control by the teacher using GhostWriter.

The teacher's control is a single slider which controls the dimension of "scariness" – that is, how frightening the music is supposed to sound. A minimum setting yields a rather banal, wistful style, while the maximum yields discordant, suspenseful music.

Here, I focus on the harmonic aspects of the system, though the melodic and rhythmic sides are also important in achieving presence.

The changing nature of Herman's music is achieved, as in the work of Bernard Herrmann, by breaking the rules of standard harmony, as elicited from Schoenberg (1978), in certain controlled ways. This is performed within a global structure known as a *form*, which determines how sections of music will relate to each other. For example, a form of AABA states that a passage (A) will repeat, be followed by a different passage (B) and then repeat again; a form of AA'A" describes the successive variation of an initial passage. Given the form, which specifies where the musical stopping points are, standard rules of harmony can be applied. As is the case in many music generation systems, choices are resolved randomly – but Herman is unusual in that the range of choices is not uniformly distributed.

Options are chosen by roulette-wheel selection from the possibilities available, but as the "scariness" level rises, the probability distributions change. Specifically, more options are made available, and increasingly likely, which break certain harmonic rules. One simple correspondence is the introduction of non-harmonic, dissonant notes into the harmony as scariness increases; another is the transition from major keys into minor ones, commonly associated with impending doom (or at least sadness) in Western music.

Also affected by the change from "neutral" to "scary" are the sounds used to play the music. In neutral mode, a gentle flute tone, in a low register, is used; as the music becomes more frightening, a higher-pitched orchestral string sound is substituted.

The important thing to understand here is that Herman's control slider encodes an approximation to the relationship between the music and the affect response experienced (or at least reported) by the user. This makes it different from any comparable system reported in the literature, and constitutes a step towards capturing musical intention in an AI system, as I shall argue in §5.

Herman's harmony and melody generators, and its rhythm generator, are described in full detail by Stapleford (1998) and de Quincey (1998), respectively.

Results

Fully formal experiments verifying the success of Herman's approximation to musical affect are still current. However, informal (but nevertheless careful) experiments with a body of about 40 adults suggest that the work is indeed successful, with respondents unanimously reporting responses as predicted.

4.2 A neural network which captures musical tension

The System

Melo (1998) reports on a connectionist system, consisting of two coupled neural networks, which models musical tension in the sense outlined above, as it varies though the duration of a piece. The two networks, working at the bar and chord levels, were trained via back-propagation, with input from the score of an extended piece of Classical music (the last movement of Prokofiev's 1st Symphony), and with the response of 10 listeners who were asked to indicate musical tension by pushing against a sprung rubber wheel. Musical tension was described to them as

> "the uneasy – as opposed to relaxed – sound of the music. Tension at a given moment reflects how 'unfinished' would the piece sound if it stopped in that precise moment."

To try to focus the attention of the respondents on *harmonic* tension, as opposed to any other kind, the music was played back on a synthesised piano, thus removing any tension due to variations in instrumentation.

There was a strong correlation between the responses of 8 of the subjects, with a ninth showing some correlation, but saturating the controller (*i.e.,* pushing hard all the time). The two outliers were omitted from the study.

The data was presented to the network as a matrix whose columns were the notes of the keyboard (Boolean) plus a representation of the reported tension, (real, between 0 and 1). The rows of the matrix represented successive beats, and these were presented to the networks as time series, thus removing the open question in the connectionist work of how much memory is presented to the network – here, this is fixed. The output of the network at any given time was trained to be the next chord in the piece.

Two networks were trained, one with bar data (the harmonic content of each bar) and one with note data (with a memory of just three beats). The bar network's output was fed to the note network, also. Once the networks were trained, a seed was used to start a sequence of musical output (both random, and from composed music), and then feedback was used to produce a continuous stream of musical progressions, the tension input being determined by the user.

Results

The results of the experiment were mixed. In the first instance, because pure feedback without perturbation is being used, the network tends to settle down into a very simple (though harmonically correct and Prokofiev-like) repeating sequence. However, even though this is the case, variation of the tension input does cause variaton of the output in ways which may objectively be said to increase or decrease the tension.

A formal verification of the results by experiment is still current, but the working hypothesis is that musical tension can indeed be captured in this approach. Even though this was a small scale experiment, which cannot be called conclusive, it suggests that there is considerable scope for the study of human musical comprehension by means of this and similar approaches.

5 Discussion

How, then, do the two systems I describe in §4 lead to an answer to to my original question: "what's missing" from automated musical composition systems?

Clearly, there is a long way to go before we can address the mystical nature of the intended ebb and flow demanded by listeners. But these two case studies show that it is possible at least to approximate aspects of music other than its stylistic correctness rules, as (more or less) captured by all of the existing systems reviewed in §3.

Specifically, Herman demonstrates that (some of) the musical qualities which give rise not just to affective response, but to affective response *of a particular kind*, can be captured in a rule-based system. Even if its output is not able to engender actual fear (and indeed, this would not be desirable in context), its listeners are able to identify that its output *suggests* fear, which is what is needed.

Melo's system shows that, so long as agreement on the intention behind the music can be achieved, non-symbolic systems can learn to reproduce it – in spite of the fact that the method is not, in this simple form, capable of producing high-quality harmony!

These two ideas, while not adequate to explain musical intention in themselves, are important components of a larger system which might be able to do so. It is now necessary, therefore, to carry this work forward, and to broaden it to the other aspects of musical intention not covered here.

6 Conclusion

In this paper, I have suggested that there is a missing ingredient, *musical intention*, in the existing AI systems which attempt to solve the problem of musical harmonisation. While not proposing an answer, I have indicated how we may begin to address the problem, both in symbolic, knowledge-rich systems, and in non-symbolic, learning-based systems.

This exercise is not a luxury; it is not merely the icing on the cake of musical composition. Without it, the fundamental point of music-making is wildly missed, and AI's efforts in the direction of musical composition become pointless. It is to be hoped that more efforts in this direction will give rise to some more concrete answers in the near future, so that the efforts of researchers in the field of intelligent music systems may bear fruit of higher and more obvious musical quality than hitherto.

References

Angluin, D. (1980). Finding patterns common to a set of strings. *Journal of Computing Systems Science 21*, 46–62.

Barucha, J. (1987). Music cognition and perceptual facilitation: a connectionist framework. *Music Perception 5*, 1–30.

Bel, B. and J. Kippen (1992). Bol processor grammars. In O. Laske, M. Balaban, and K. Ebcioglu (Eds.), *Understanding Music with AI – Perspectives on Music Cognition*, pp. 366–401. Cambridge, MA: MIT Press.

Biles, J. (1994). Genjam: A genetic algorithm for generating jazz solos. In *Proceedings of the International Computer Music Conference*.

de Quincey, A. (1998). Herman. Master's thesis, Division of Informatics, University of Edinburgh.

Ebcioglu, K. (1988). An expert system for harmonizing four-part chorales. *Computer Music Journal 12*.

Ebcioglu, K. (1990). An expert system for harmonizing chorales in the style of J. S. Bach. *Journal of Logic Programming 8*.

Gang, D., D. Lehmann, and N. Wagner (1997). Harmonizing melodies in real time: the connectionist approach. In P. R. Cook (Ed.), *Proceedings of ICMC'97*.

Krumhansl, C. (1997). Musical tension: Cognitive, motional and emotional aspects. In *Proceedings of the 3rd Triennial ESCOM Conference*. European Society for the Cognitive Sciences of Music.

Lévi-Strauss, C. (1964). *The Raw and the Cooked*. London: Jonathan Cape. Volume 1 of Introduction to a Science of Mythology.

Longuet-Higgins, H. C. (1962). Letter to a musical friend. *The musical review 23*, 244–8,271–80.

Melo, A. F. (1998). A connectionist model of tension in chord progressions. Master's thesis, Division of Informatics, University of Edinburgh.

Meyer, L. (1956). *Emotion and Meaning in Music.* University of Chicago Press.

Ovans, R. (1992). Efficient music composition via consistency techniques. Technical Report CSS-IS TR 92-02, Centre for Systems Science, Simon Fraser University, Burnaby, B.C., Canada.

Pachet, F. and P. Roy (1998). Reifying chords in automatic harmonization. In F. Pachet and P. Codognet (Eds.), *Proceedings of the ECAI'98 Workshop on Constraints for Artistic Applications.*

Papadopoulos, G. and G. A. Wiggins (1998). A genetic algorithm for the generation of jazz melodies. In *Proceedings of STeP'98*, Jyväskylä, Finland.

Papadopoulos, G. and G. A. Wiggins (1999). AI methods for algorithmic composition: A survey, a critical view and future prospects. Technical report, Division of Informatics, University of Edinburgh. Submitted to AISB'99.

Phon-Amnuaisuk, S., A. Tuson, and G. A. Wiggins (1999). Evolving musical harmonisation. In *Proceedings of ICANNGA'99*, Portorož, Slovenia.

Ponsford, D., G. A. Wiggins, and C. S. Mellish (FC). Statistical learning of harmonic movement. *Journal of New Music Research*. Forthcoming. Also available as Research Paper 874, from the Division of Informatics, School of Artificial Intelligence, University of Edinburgh.

Robertson, J., A. de Quincey, T. Stapleford, and G. A. Wiggins (1998). Real-time music generation for a virtual environment. In F. Nack (Ed.), *Proceedings of the ECAI'98 Workshop on AI/Alife and Entertainment*, Brighton, England. Also available as a Research Paper from the Department of Artificial Intelligence, University of Edinburgh.

Schoenberg, A. (1978). *Theory of Harmony* (3 ed.). Belmont Music Publishers.

Stapleford, T. (1998). The harmony, melody, and form of herman, a real-time music generation system. Master's thesis, Division of Informatics, University of Edinburgh.

Steedman, M. J. (1994). The well-tempered computer. *Transactions of the Royal Society of London, Series A 349*, 115–131.

Steedman, M. J. (1996). The blues and the abstract truth: Music and mental models. In *Mental Models In Cognitive Science*, pp. 305–318. Mahwah, NJ: Erlbaum.

Steedman, M. J. (1999). Categorial grammar. In *The MIT Encyclopedia of Cognitive Sciences.* MIT Press.

Sundberg, J. and B. Lindblom (1991). Generative theories for describing musical structure. In P. Howell, R. West, and I. Cross (Eds.), *Representing Musical Structure.* Academic Press.

Taylor, E. (1996). *The* **AB** *guide to music Theory.* London: The Associated Board of the Royal Schools of Music.

Wiggins, G. A. (1998a). Music, syntax, and the meaning of "meaning". In *Proceedings of the First Symposium on Music and Computers*, Corfu, Greece.

Wiggins, G. A. (1998b). The use of constraint systems for musical composition. In F. Pachet and P. Codognet (Eds.), *Proceedings of the ECAI'98 Workshop on Constraints and Artistic Applications*, Brighton, England. Also available as a Research Paper from the Department of Artificial Intelligence, University of Edinburgh.