



Faculty Publications

2008-01-01

Automatic Composition of Themed Mood Pieces

Heather Chan
heather.y.chan@gmail.com

Dan A. Ventura
ventura@cs.byu.edu

Follow this and additional works at: <https://scholarsarchive.byu.edu/facpub>



Part of the [Computer Sciences Commons](#)

Original Publication Citation

Heather Chan and Dan Ventura, "Automatic Composition of Themed Mood Pieces", Proceedings of the International Joint Workshop on Computational Creativity, pp. 19-115, 28.

BYU ScholarsArchive Citation

Chan, Heather and Ventura, Dan A., "Automatic Composition of Themed Mood Pieces" (2008). *Faculty Publications*. 908.

<https://scholarsarchive.byu.edu/facpub/908>

This Peer-Reviewed Article is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Faculty Publications by an authorized administrator of BYU ScholarsArchive. For more information, please contact ellen_amatangelo@byu.edu.

Automatic Composition of Themed Mood Pieces

Heather Chan and Dan Ventura

Department of Computer Science
Brigham Young University
Provo, Utah, USA
heatherchan@byu.net, ventura@cs.byu.edu

Abstract. Musical harmonization of a given melody is a nontrivial problem; slight variations in instrumentation, voicing, texture, and bass rhythm can lead to significant differences in the mood of the resulting piece. This study explores the possibility of automatic musical composition by using machine learning and statistical natural language processing to tailor a piece to a particular mood using an existing melody.

Key words: Composition, Harmonization, Creativity, Mood

1 Introduction

There are many examples of thematic musical works in which the composer takes a single melody and reworks it in different ways to produce a variety of related material. This practice is particularly useful in soundtracks for programmatic works such as films, TV series, and video games. By referring back to and reusing the same melodic elements, the composer is able to tie together several otherwise unrelated pieces so that listeners understand and associate with a specific storyline and set of characters. Examples of such works include the Lord of the Rings movie trilogy, the Simpsons TV series, and the Zelda game series.

The reworking of an existing theme is a skill subset of composition that demands creativity from the composer, yet can be viewed as a simpler task than composition from scratch. We are interested in the ability of software to emulate this compositional practice in order to produce an extensive set of musical material for use in the aforementioned genres. In particular, we wish to examine the feasibility of tailoring individual pieces to specific moods based on user request.

Approaches to automatic harmonization include n -gram statistical learning for learning musical grammars [1], genetic algorithms for generating four-part harmony [2], and hidden Markov models for chorale harmonization [3]. We have chosen to concentrate on statistical methods, which can apply one of at least four basic methods of music generation: random walk, HMM with Viterbi decoding, stochastic sampling and pattern-based sampling [4]. Chuan and Chew and Whorley *et al.* have investigated the automatic generation of style-specific accompaniment for melodies [5, 6]. This is very similar to what we hope to accomplish in that their systems address melodies that have already been composed

and attempt to provide a well-written accompaniment for each of them. However, while their systems concentrate on musical style as the main determining factor of their harmonizations, we wish to focus on mood and emotional impact. Considering the motivational/methodological ontology suggested by Pearce *et al.* [7], we suggest that our proposed system be considered as an approach to algorithmic composition with perhaps some shading towards the design of a compositional tool.

We seek to design a system that takes an existing melody as input and produces as output an automatically composed piece of music based on user-specified parameters governing the overall mood or feel of the piece. We limit our scope to melodies in regular meters that do not contain non-chord tones on strong beats and that do not call for key modulation in the harmony.

Many systems for automatic musical composition have been met with criticism for being either too biased by the tastes of the creators or, conversely, too broad or arbitrary in their definitions of music to be able to meet the musical tastes of the audience. We aim to develop a system that demonstrates a higher level of creativity than these by creating independently and creating something of value.

2 System Design

Our design requires four functions: a *melody filter* to select the melody notes we will use to harmonize; a *chord progression generator* to produce a suitable chord progression for the melody subset; a *harmonization planner* to select the instrumentation and harmonization technique to be used based on the mood; and a *composer* to build the piece from the melody, the generated chord progression, and the chosen instrumentation and harmonization technique (see Fig. 1).

2.1 Melody Filter

For now, we implement the melody filter as simply as possible. We specify a sampling interval based on measure length over which a single chord will persist, based on strong beats, and we consider only the first melody note in each interval as significant in chord assignment. An example is shown in Figure 2.

2.2 Chord Progression Generator

The chord progression generator requires a more complex approach. To harmonize a given melody, one must consider not only the melody notes but also which chords are compatible with them and which of those chords fit together well. We can treat the melody notes as observed events and the underlying chord progression as a hidden state sequence, modeling the inference problem as a simple hidden Markov model (HMM) (see Figure 3). Here, we must determine the chord progression from the melody notes, similar to the process of Roman numeral analysis used by music theory students. The first-order HMM depicted

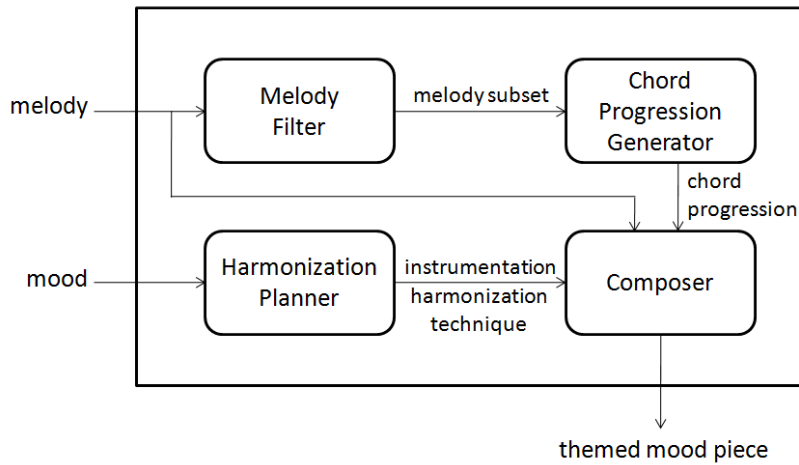


Fig. 1. Block diagram of system design. Given an input melody, the *melody filter* selects the melody notes that will be used to harmonize. Those notes are passed to the *chord progression generator* which produces a suitable chord progression for the melody subset. The *harmonization planner* selects the instrumentation and harmonization technique to be used based on the desired mood. Finally, the *composer* builds the piece from the melody, the generated chord progression, and the chosen instrumentation and harmonization technique.

in Figure 3 should be sufficient for modeling well-formed chord progressions, as classical music theory suggests that the extent to which each successive chord is pleasing to the ear is heavily conditioned on its relationship to the immediately preceding chord.

The parameters of the model represent two conditional probability distributions: the next chord given the preceding chord $P(c_i|c_{i-1})$ and the observed melody note given the current chord $P(m_i|c_i)$. We expect the former to demonstrate that the dominant almost always resolves to the tonic, and occasionally to the submediant; the latter should reflect the low probability of observing the second scale degree in the melody line during a tonic chord. These distributions can be modeled using statistics gathered from MIDI datasets.

Given the melody, we can use statistical inference to sample the distribution $P(c_i|m_i)$, giving us, for example, a maximum likelihood estimate of the generating chord progression; we have chosen the Viterbi algorithm for now to accomplish this¹. The previous example is continued in Figure 4. When a suitable chord progression has been determined, we can build harmonizations using common accompanimental figures.

¹ We limit our study to diatonic chords with no key modulation.




Fig. 2. *Melody filtering.* “Twinkle, Twinkle Little Star” is sampled in half-measure intervals and only the sampled notes are considered significant for determining chord progression.




Fig. 3. *Graphical model used for chordal inference and progression generation.* The nodes labeled c_1, c_2, \dots, c_n represent the hidden chords, and the nodes labeled m_1, m_2, \dots, m_n represent the observed (sampled) melody notes.

2.3 Harmonization Planner

We view the implementation of the harmonization planner as a machine learning problem. Initially, we have limited our choice of instrumentation to piano accompaniment, string accompaniment, harp accompaniment, and electric guitar accompaniment. We would like to be able to specify a mood to the system such as “joyous”, “urgent”, “tragic”, etc. and have the system compose a suitable piece for that mood. In order to learn the appropriate parameters, we need to rely on human feedback, as we have no way of determining the effective mood of a piece without referring to popular opinion.

We will investigate two different approaches to learning the relationship between instrumentation, harmonization, and mood. The first method will involve generating several pieces of music using different instrumentations and harmonizations as input and asking listeners to categorize the resulting pieces into different classes of moods; we can then cluster the examples and select a nearest neighbor example of instrumentation and harmonization when a user requests a specific mood. This will require a way to define a distance measure between



Fig. 4. *Chord progression.* Sampling the distribution $P(c_i|m_i)$ results in a possible generating chord sequence for “Twinkle, Twinkle Little Star”.



Fig. 5. *Composition.* An arpeggiated harmony for “Twinkle, Twinkle Little Star”.

pieces that will yield useful clusters, possibly through a trial-and-error weighting of the inputs and manual inspection of the resulting clusters.

The second method will involve having the user specify a mood, generating several pieces of music as before, and having the user choose the n best pieces that come the closest to the mood effect that the user had in mind; the input would be the specified mood, and the outputs would be the specific instrumentation and harmonization technique used. This is a single-input multiple-output learning problem that may be difficult to learn but may be amenable to solution by neural networks or by application of inverse methods.

2.4 Composer

Once the chord progression, instrumentation, and harmonization technique have been determined, the final piece can be written. This is a simple matter of setting the voices according to the selected instrumentation, expanding the chords in the chord progression using the chosen harmonization technique, and combining the resulting accompaniment with the original melody. The concluding result of our example is shown in Figure 5.

3 Assessment of Creativity

The artist’s ambition is not to discover a “correct” answer, but a “creative” one. We have therefore decided to approach the evaluation of this system not as a measure of accuracy, but as one of creativity.

Recently it has been suggested that for a computational system to be considered creative, it must be perceived as possessing *skill*, *appreciation*, and *imagination* [8]. We can describe the system as skillful if it exhibits knowledge of musical behavior and can make intelligent decisions in composition that lead to a well-formed piece of music. We have provided our system with statistical information about accepted progressions and behavior via an HMM, and we are implementing compositional techniques for harmonization including a few different forms of chordal accompaniments. We thus argue that this system is skillful.

We can describe the system as appreciative if it consistently produces something of value and can evaluate itself to some extent in order to adapt its work to the user's tastes. We are currently working on using machine learning algorithms and user feedback to discover associations between various instrumentations and harmonization techniques and the moods they tend to produce. We hope to amass enough data from user feedback to make this possible; if we succeed, we will be able to argue that this system is appreciative.

We can describe the system as imaginative if it creates new material that demonstrates some level of independence both from its creator's designs and from works by other composers. We have introduced a bias towards particular datasets by relying on statistical methods to ensure that the system produces plausible chord progressions; we have also introduced a bias towards the creator's compositional style because the system draws only on instrumentations and harmonization techniques that are programmed. However, these influences are acceptable to some degree since all composers are influenced both by their mentors and by other composers whose works they have studied. In addition, this system has two creative abilities that help offset bias: it is able to generate chord progressions that have never been seen before, and it selects combinations of instrumentations and harmonization techniques based on knowledge of user feedback, which is independent both of the creator's tastes and of the tastes of other composers. We therefore describe this system as somewhat imaginative.

Although it is difficult to quantify creativity, we can survey audiences and obtain numerical ratings for each of these characteristics. Questions might include: "On a scale of 1 to 10, how much does this sound like real music?" "On a scale of 1 to 10, how closely does this match the mood you were thinking of?" "On a scale of 1 to 10, how much does this sound like something you have heard before?" (skill, appreciation, and imagination, respectively). In addition to being informative on their own, these ratings might be useful data for training the system to self-evaluate its creativity and increase its level of appreciation.

Our system is skillful and somewhat imaginative, but still lacks appreciation for the value of what it can create. We hope to remedy this by exposing the system to more user feedback in the future in order to help it learn what is of value and what is not.

4 Discussion

We have proposed a system for automatic composition of mood pieces based on an existing melody and have presented our approach to constructing such a system. Currently, the system design is limited to diatonic melodies that do not modulate in key, but we plan to accommodate greater harmonic complexity and key modulation in the input melodies in the future, expanding the system's harmonic vocabulary to include borrowed and altered chords and investigating the use of a key-finding algorithm to detect key modulation.

Of course, greater harmonic complexity will significantly increase the size of the joint space, making our inference task computationally challenging. To address this issue, we plan to incorporate stochastic inference methods to replace the use of the Viterbi algorithm. As an additional benefit, the stochastic element will allow for nondeterminism in the selection of an appropriate chord progression, enhancing the system's creative ability.

We are currently limited in our choice of instrumentations and harmonization techniques, and we would like to incorporate percussion rhythms and effects, as these are often crucial in heightening the emotional intensity of a piece. If possible, we would also like to implement other tools of variation such as change of register, change of meter, change of mode, and melody augmentation/diminution.

Finally, as it is difficult to gather data for the mood-learning portion of this problem (since it must be obtained from a human audience), we require an efficient method of obtaining significant amounts of data on the correlations between instrumentation, harmonization, and mood.

References

1. Ponsford, D., Wiggins, G., Mellish, C.: Statistical learning of harmonic movement. *Journal of New Music Research* **28**(2) (1998) 150–177
2. Phon-Amnuaisuk, Wiggins, G.: The four-part harmonisation problem: a comparison between genetic algorithms and a rule-based system. In: *Proceedings of the AISB99 Symposium on Musical Creativity*. (1999)
3. Allan, M., Williams, C.K.I.: Harmonising chorales by probabilistic inference. In: *Advances in Neural Information Processing Systems* 17. (2005) 25–32
4. Conklin, D.: Music generation from statistical models. In: *Proc. AISB Symposium on Artificial Intelligence and Creativity in the Arts and Sciences*. (2003) 30–35
5. Chuan, C., Chew, E.: A hybrid system for automatic generation of style-specific accompaniment. In: *Proc. Int. Joint Workshop on Computational Creativity*. (2007) 57–64
6. Whorley, R.P., Wiggins, G.A., Pearce, M.T.: Systematic evaluation and improvement of statistical models of harmony. In: *Proceedings of the International Joint Workshop on Computational Creativity*. (2007) 81–88
7. Pearce, M., Meredith, D., Wiggins, G.: Motivations and methodologies for automation of the compositional process. *Musicae Scientiae* **6**(2) (2002) 119–147
8. Colton, S.: Creativity vs. the perception of creativity in computational systems. In: *Creative Intelligent Systems: Papers from the AAAI Spring Symposium*, Technical Report SS-08-03, AAAI Press (2008) 14–20