



## **Math & Music: The Perfect Match**

Elaine Chew

### **► To cite this version:**

| Elaine Chew. Math & Music: The Perfect Match. 2008, <10.1287/orms.2008.03.12>. <hal-03687047>

**HAL Id: hal-03687047**

**<https://hal.science/hal-03687047v1>**

Submitted on 3 Jun 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

## Music and Operations Research - The Perfect Match?

By Elaine Chew

**Operations research (OR) is a field that prides itself in its versatility in the mathematical modeling of complex problems to find optimal or feasible solutions. It should come as no surprise that OR has much to offer in terms of solving problems in music composition, analysis, and performance. Operations researchers have tackled problems in areas ranging from airline yield management and computational finance, to computational biology and radiation oncology, to computational linguistics, and in eclectic applications such as diet formulation, sports strategies and chicken dominance behavior, so why not music?**

I am an operations researcher and a musician. Over the past ten years, I have devoted much time to exploring, and building a career at, the interface between music and operations research. At the University of Southern California, I founded the Music Computation and Cognition (MuCoaCo) Laboratory<sup>1</sup>, and direct research in music and computing. To bring music applications to the attention of the OR community, I have organized invited clusters at INFORMS meetings: 'OR in the Arts: Applications in Music'<sup>2</sup> at the 2003 INFORMS meeting in Atlanta; and, 'Music, Computation and Artificial Intelligence'<sup>3</sup> at the 2005 INFORMS Computing Society (ICS) meeting in Annapolis. The summer 2006 issue (volume 18, issue 3<sup>4</sup>) of the INFORMS Journal on Computing (JoC) featured a special cluster of papers on Computation in Music that I edited with Roger Dannenberg, Joel Sokol, and Mark Steedman.

The goal of the present article is to provide a broader perspectives on the growth and maturation of music and computing as a discipline, resources to learn more about the field, and some ways in which operations research impacts, and can influence, this rapidly expanding field. My objective is to show that music – its analysis, composition, and performance – presents rich application areas for OR techniques. Any attempt to give a comprehensive overview of the field in a few pages must necessarily be doomed to failure, due to the vastness of the domain. I shall give selected examples, focusing on a few research projects at the MuCoaCo Laboratory and some related work, of musical problems that can be framed and solved mathematically and computationally. Many of the techniques will be familiar to the OR community, and some will borrow from other computing fields. These examples will feature the modeling of music itself, the analysis and description of its structures, and the manipulation of these structures in composition and performance.

### Why Model Music Mathematically or Computationally?

The widespread access to digitally encoded music, from desktop browsers to handheld devices, has given rise to the need for mathematical and computational techniques for generating, manipulating, processing, storing, and retrieving digital music. For example, Shazam<sup>5</sup> provides a service for users to dial a number, hold their cell phones to a music source, and receive a text message with the name of the song and recording for future purchase. Pandora<sup>6</sup> recommends songs with music qualities similar to the user's current favorite; much of the technology still depends on human annotation of existing soundtracks. A possible future application could

be a musical Google, whereby one could retrieve music files by humming a melody or by providing an audio sample.

Beyond such industrial interests in mathematical and computational modeling of music, such approaches to solving problems in music analysis and in music making have their own scientific and intellectual merit as a natural progression in the evolution of the disciplines of musicology, music theory, performance, composition, and improvisation. Formal models of human capabilities in creating, analyzing, and reproducing music serve to further knowledge in human perception and cognition, and advance the state of the art in psychology and neuroscience. By modeling music making and analysis, we gain a deeper understanding of the levels and kinds of human creativity engaged by these activities.

The rich array of problems in music analysis, generation (composition/improvisation), and rendering (expressive performance) present new and familiar challenges to mathematical and computational modeling and analytical techniques, the bread and butter of operations researchers, in a creative and vast domain.

### Publication Venues

While the use of computers in music can be traced back to the beginnings of computers themselves, the use of mathematics and computing to mimic human intelligence and decision-making when appreciating or making music has gathered momentum only in recent years. The escalating activity in mathematical and computational research in music has led to the birth of new societies, conferences and journals. For example, I am a founding member and secretary of the Society for Mathematics and Computation in Music<sup>7</sup>, formed in 2006 by an international group of researchers. The society unveiled the inaugural issue of its flagship publication, the Journal of Mathematics and Music<sup>8</sup> at its first meeting in May of 2007. A special issue on Computation will appear in summer 2008, edited by Alfred Cramer, Christopher Raphael, and myself. More traditional archival publication venues include the Journal of New Music Research<sup>9</sup>, the Computer Music Journal<sup>10</sup>, and Computing in Musicology<sup>11</sup>.



(left) Journal of the Society of Mathematics and Computation in Music. (right) Editors-in-chief, Peck and Noll, and society members amongst speakers of the Mathematical Techniques in Musical Analysis sessions at the Joint Mathematics Meeting, New Orleans, January 6, 2007. Left-to-right: Franck Jedrzejewski, William Sethares, Godfried Toussaint, Robert Peck, Thomas Noll, Elaine Chew, Rachel Hall, Julian Hook, Ching-Hua Chuan, Kathryn Elder.

<sup>1</sup> MuCoaCo Laboratory: [www-ref.usc.edu/~mucoaco](http://www-ref.usc.edu/~mucoaco)

<sup>2</sup> INFORMS OR in the Arts: [www-ref.usc.edu/~echew/INFORMS/cluster.html](http://www-ref.usc.edu/~echew/INFORMS/cluster.html)

<sup>3</sup> ICS Music, Computation, AI: [www-ref.usc.edu/~echew/INFORMS/ics2005.html](http://www-ref.usc.edu/~echew/INFORMS/ics2005.html)

<sup>4</sup> INFORMS JoC vol. 18 no. 3: [joc.journal.informs.org/content/vol18/issue3](http://joc.journal.informs.org/content/vol18/issue3)

<sup>5</sup> Shazam: [www.shazam.com](http://www.shazam.com)

<sup>6</sup> Pandora: [www.pandora.com](http://www.pandora.com)

<sup>7</sup> Society for Mathematics and Computation in Music: [www.smcmm-net.info](http://www.smcmm-net.info)

<sup>8</sup> Journal of Mathematics and Music: [www.tandf.co.uk/journals/titles/17459737.asp](http://www.tandf.co.uk/journals/titles/17459737.asp)

<sup>9</sup> Journal of New Music Research: [www.tandf.co.uk/journals/titles/09298215.asp](http://www.tandf.co.uk/journals/titles/09298215.asp)

<sup>10</sup> Computer Music Journal: [204.151.38.11/cmj](http://204.151.38.11/cmj)

<sup>11</sup> Computing in Musicology: [www.ccarb.org/publications/books/cm](http://www.ccarb.org/publications/books/cm)





configuration to exploit its continuous interior space, and uses the convexity of tonal structures, such as chords and keys, to define spatial representations of these objects in the interior.


Automating Music Analysis

A discussion of music representation naturally leads mathematical formulations of problems in music analysis. At a crude level, tonal compositions can be organized into sections, each section consists of regions in particular keys that determine the predominant tones that are sounded in that region, each key region consists of segments of chord sequences, and chord comprises of tone clustered in or over time. Music and computing researchers have sought ways to automatically abstract these structures from music data, using both discrete and probabilistic methods.

In tonal music, the key refers to the predominant pitch set used in the music, and is identified by what is perceived to be the most stable tone. An early key finding algorithm by Longuet-Higgins and Steedman (1971) used shape matching on the harmonic network to determine the key. Inspired by Longuet-Higgins and Steedman’s method and by interior point approaches to linear optimization, the center of effect generator (CEG) algorithm by Chew (2000) uses the three-dimensional configuration of the harmonic network, and the interior space to track the evolving tonal context. In the CEG method, a piece of music or its melody generates a sequence of centers of effect that traces a path in the interior of the array of pitches. The key at any given point in time is computed by a nearest neighbor search for the closest major or minor key representation on the respective major/minor key helices. By moving from the lattice to the interior space, the model becomes more robust to noise, and finds the key in fewer note events.


E	B	F#	C#	G#	D#	A#	E#	B#
C	G	D	A	E	B	F#	C#	G#
Ab	Eb	Bb	F	C	G	D	A	E
Fb	Cb	Gb	Db	Ab	Eb	Bb	F	C
Dbb	Abb	Ebb	Bbb	Fb	Cb	Gb	Db	Ab

The Harmonic network and Longuet-Higgins and Steedman’s major (red) and minor (blue) key templates for their key finding algorithm.



Pitch class (silver), major key (red), and minor key (blue) helices in the spiral array model. The squiggly green line shows a center of effect path produced by a hypothetical segment of music, and the green box indicates the closest key. The translucent blue shape covers the convex hull of the pitches in the closest key, which happens to be a minor key.

The spiral array model and its associated tonal analysis algorithms have been implemented in the MuSA.RT (Music on the Spiral Array . Real-Time) system for interactive tonal analysis and visualization (Chew & François, 2005). Any operations researcher who has implemented a mathematical algorithm for computational use can attest that it is one



Composer Lisa Bielawa tries out the MuSA.RT music analysis and visualization system following Chew and François’ presentation at the Radcliffe Gymnasium, January 16, 2008. Photo by E. Chew.

thing to design an algorithm of low computational complexity, and another to implement it so that it runs in real time. MuSA.RT is built using François’ Software Architecture for Immersipresence framework (2004), which allows for the efficient processing of multiple concurrent data streams. MuSA.RT has been featured in numerous presentations, including *The Mathematics in Music* concert-conversation, which made its debut in Los Angeles, in Victoria, B.C., and in Singapore in 2007. Venues in 2008 include North Carolina State University and the Massachusetts Institute of Technology in the Eastern United States.

Apart from the pitch structures described in the previous paragraphs, music also possesses time structures. When listening to music, humans are quickly able to pick out the beat, and to tap along with it. The patterns of long and short durations along the grid of the beat produce rhythm. Periodicity in the accent patterns lead to the establishing of meter. Using the mathematical model for metrical analysis, developed at the Multimedia Laboratory headed by Guerino Mazzola (now at the University of Minnesota), and described and advanced by Anja Volk (2002), Chew, Volk, and Lee (2005) proposed a method for classifying dance music by analyzing their metrical patterns, providing an alternative to the inter-onset-interval distribution and





MIMI's international debut at the Musical Instrument Museum, Berlin, with the Seiler piano – the first MIDI grand piano, built in the 50s for TV production, May 2007. Photo by E. Chew.

autocorrelation methods proposed by Dixon, Pampalk, and Widmer (2003).

Apart from the issues of determining tonal and rhythmic contexts, which can be thought of as finding vertical structures, there are equally interesting and challenging questions with regard to the determining of horizontal structures, such as melody and voice. A number of models have been proposed for finding motivic patterns, an example of a motive is the opening four notes of Beethoven's Fifth, and for separating voices, independent melodic threads superimposed in time. Their sequential nature means that techniques inspired by DNA sequence analysis and tree structure approaches lend themselves readily to solving problems of motivic pattern discovery – see, for example, Conklin & Anagnostopoulou (2006) and Lartillot (2005). In the same strain of computational biology-inspired approaches, we proposed a contig mapping approach to voice separation by first fragmenting, then assembling, the voices that make up a polyphonic (multi-voice) composition (Chew & Wu, 2004).

### Music Composition and Improvisation

One of the first applications of music and computing that comes to mind is the use of mathematical and computational models for generating music. Music composition and improvisation can often be conceived of, and modeled as, the exploration of a solution space bounded by constraints imposed by the composer or improviser. Truchet & Codognot (2004) modeled a number of problems in computer-assisted composition as constraint satisfaction problems. A few composers have also discovered operations research techniques: Tom Johnson (2004) poses composition problems to mathematicians / operations researchers, for example tracing a Hamiltonian circuit through a network of  $n$ -note chords connected to other  $n$ -note chords that share  $k$  pitches, and Daniel Schell (2002) explores the ideas of optimality in harmonic progressions.

Taking advantage of the fact that music bears similarities with language, Steedman (1996) proposed a formal grammar for generating jazz chord sequences. Computational linguistics approaches typically require large corpora of annotated data for learning, which poses numerous challenges for music data. In efforts spearheaded by Ching-Hua Chuan and Reid Swanson, we are seeking ways to generate style-specific accompaniment given only a few examples (Chuan & Chew, 2007), and to segment melodies using unsupervised techniques so as to assist in computational creativity projects in music (Swanson, Chew & Gordon, 2008).

Most systems for generating music originate in some Markov model, as exemplified by Pachet's *Continuator* (2003), which builds a variable order Markov model from an input sequence for generating

including the region from which musical material was being sampled for recombination, ten seconds' lead time before the sounding of the new material, and an equal amount of time to review the improvisation, and to plan future strategies during the performance. MIMI received its international debut at the Musical Instrument Museum in Berlin in May, 2007, at the International Conference on Mathematics and Computation in Music.

### Expressive Music Performance

While the music score specifies the notes to be played, and to some degree how they should be played, it contains ambiguities the resolution for which can be influenced by a performer. Interpretation in music performance has numerous analogies in speech, where the expression with which a text is spoken, the grouping of words in a sentence, and strategic pauses can affect the meaning of, and add nuance to, the words. The devices by which a performer injects expression into a score include timing deviations, loudness modulations, and articulation (the shortening and lengthening of note durations). A performance can be viewed as a set of decisions regarding the employment of these expressive strategies, constrained by the structure of the notes in the score. Researchers have worked to understand and model the principles of expressive strategies, for example, using machine-learning tools (Widmer, 2001), or through direct observation of empirical data (Palmer, 1996, and Repp, 1998, 1999).


In the Expression Synthesis Project (Chew et. al., 2005, 2006) we employ an analysis-by-synthesis approach to study expressive performance. Taking the cue from motion capture for the generation of realistic animation, ESP uses the metaphor of driving for expressive performance, and a driving (wheel and pedals) interface to map car motion to tempo (speed) and amplitude in the rendering of an expressionless piece. The road is designed to guide expression: bends in the road encourage slowdowns, while straight sections promote speedups. Buttons on the wheel allow the user to change the articulation, by shortening or lengthening the notes. A virtual radius mapping strategy (Liu et. al.) ensures tempo smoothness.

A goal of ESP is to make high level decision making in expressive performance widely accessible to both experts and novices alike. The system was exhibited at the University of Southern California Festival 125 Pavilion over three days in October, 2005. It was showcased again at the National University of Singapore Techno-Arts Festival in March 2007.

One cannot discuss expression without considering its effect on the emotion perceived or experienced by the listener. In Parke, Chew, and Kyriakakis (2007ab), we present quantitative modeling, visualization, and regression analyses of the emotion perceived in film



ESP at the Festival 125 Pavilion at USC. The pavilion housed exhibits, and was visited by 10,000 visitors over three days, October 6-8, 2005. Photos by E. Chew.



E. Chew and X. Wu, 2004. "Separating Voices in Polyphonic Music: A Contig Mapping Approach," in Uffe K. Wiil (ed.): *Computer Music Modeling and Retrieval: Second Intl. Symposium, CMMR 2004, Esbjerg, Denmark, May 26-29, 2004*. Revised Papers Lecture Notes in Computer Science

with and without music, showing that perceived emotion in film with music can be reliably predicted by perceived emotion in film alone and in music alone. In Mosst's Masters thesis (2006), he models and visualizes individuals' emotion responses to music using features extracted from the music's audio. An interesting finding is that humans' responses to music differ widely, and require individualized models for accurate prediction.

## Conclusion

While many examples have been given in the areas of music representation, analysis, composition/improvisation, and expressive performance, we have only begun to scratch the surface with regard to furthering knowledge about these areas of creative human activity through mathematical and computational modeling. I hope I have succeeded in demonstrating the breadth of opportunities that abound for operations researchers to make significant contributions in this fledgling field.

EC. February 6, 2008.

## Acknowledgements

This material is based in part upon work supported by the National Science Foundation under Grant No. 0347988, and Cooperative Agreement No. EEC-9529152, and by the Radcliffe Institute for Advanced Study. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation or of the Radcliffe Institute for Advanced Study.

## References

G. Assayag and S. Dubnov, 2004. "Using Factor Oracles for Machine Improvisation," *Soft Computing*, Vol. 8, No. 9, pp. 604-610.  
 E. Chew, 2000. Towards a Mathematical Model of Tonality. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.  
 E. Chew and A. R. J. François, 2005. "Interactive multi-scale visualizations of tonal evolution in MuSART Opus 2," *ACM Computers in Entertainment*, Vol. 3, No. 4, October 2005, 16 pages.  
 E. Chew, A. R. J. François, J. Liu, A. Yang, 2005. "ESP: a driving interface for Expression Synthesis," in *Proceedings of the Intl. Conference on New interfaces for musical expression*, Vancouver, Canada.  
 E. Chew, J. Liu, and A. R. J. François, 2006. "ESP: roadmaps as constructed interpretations and guides to expressive performance," in *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*, Santa Barbara, CA, pp. 137-145.  
 E. Chew, A. Volk, and C.-Y. Lee, 2005. "Dance Music Classification Using Inner Metric Analysis: a computational approach and case study using 101 Latin American Dances and National Anthems," in (eds.) B. Golden, S. Raghavan, E. Wasil: *The Next Wave in Computing, Optimization and Decision Technologies: Proceedings of the 9th INFORMS Computer Society Conference*. Kluwer.

Series (LNCS) Vol. 3310, pp. 1-20, Springer-Verlag GmbH: Berlin.  
 C.-H. Chuan and E. Chew, 2007. "A Hybrid System for Automatic Generation of Style-Specific Accompaniment," in *Proceedings of the Fourth International Joint Workshop on Computational Creativity*, Goldsmiths, University of London.  
 R. Cohn, 1997. "Neo Riemannian Operations, Parsimonious Trichords, and Their 'Tonnetz' representations," *Journal of Music Theory*, Vol. 41, No. 1, pp. 1-66.  
 D. Conklin, and C. Anagnostopoulou, 2006. "Segmental Pattern Discovery in Music," *INFORMS Journal on Computing*, Vol. 18, No. 3, pp. 285-293.  
 S. Dixon, E. Pampalk, and G. Widmer, 2003. "Classification of Dance Music by Periodicity Patterns," in *Proceedings of the Intl. Conference on Music Information Retrieval*.  
 A. Fleischer, 2002. "A Model of Metric Coherence," in *Proceedings of the 2nd Intl. Conference Understanding and Creating Music*, Caserta.  
 A. R. J. François, 2004. "A Hybrid Architectural Style for Distributed Parallel Processing of Generic Data Streams," in *Proceedings of the Intl. Conference on Software Engineering*, pp. 367-376, Edinburgh, Scotland, UK.  
 A. R. J. François, E. Chew, and C. D. Thurmond, 2007. "Visual Feedback in Performer-Machine Interaction for Musical Improvisation," in *Proceedings of the Intl. Conference on New Interfaces for Musical Expression*, New York City.  
 O. Lartillot, 2005. "Multi-Dimensional motivic pattern extraction founded on adaptive redundancy filtering," *Journal of New Music Research*, Vol. 34, No. 4, pp. 375-393.  
 T. Johnson, 2004. "Musical Questions for Mathematicians," A text presented for participants of the MaMuX seminar at IRCAM, Nov. 20.  
 J. Liu, E. Chew, and A. R. J. François, 2006. "From driving to expressive music performance: ensuring tempo smoothness," in *Proceedings of the 2006 ACM SIGCHI Intl. Conference on Advances in computer entertainment technology*, Hollywood, CA, Article No. 78.  
 H. C. Longuet-Higgins, 1962a. "Letter to a Musical Friend," *Music Review* 23, pp. 244-248.  
 H. C. Longuet-Higgins, 1962b. "Second Letter to a Musical Friend," *Music Review* 23, pp. 271-280.  
 H.C. Longuet-Higgins and M. J. Steedman, 1971. "On Interpreting Bach," in (eds.) B. Meltzer and D. Michie, *Machine Intelligence* 6, pp. 221-241. Edinburgh University Press.  
 M. Merrick, 2006. Quantitative Modeling of Emotion Perception in Music, Masters thesis, University of Southern California, Los Angeles, CA.  
 F. Pachet, 2003. "The Continuator: Musical Interaction With Style," *Journal of New Music Research*, Vol. 32, No. 3, pp. 333-341.  
 C. Palmer, 1996. "On the assignment of structure in music performance," *Music Perception*, Vol. 14, pp. 23-56.  
 R. Parke, E. Chew, C. Kyriakakis, 2007a.  
 B. H. Repp, 1998. "A microcosm of musical expression. I. Quantitative analysis of pianists' timing in the initial measures of Chopin's Etude



- in E major," *Journal of the Acoustical Society of America*, Vol. 104, No. 2, pp. 1085-1100.
- B. H. Repp, 1999. "A microcosm of musical expression: II. Quantitative analysis of pianists' dynamics in the initial measures of Chopin's Etude in E major," *Journal of the Acoustical Society of America*, Vol. 105, No. 3, pp. 1972-1988.
- D. Schell, 2002. "Optimality in musical melodies and harmonic progressions: The travelling musician," *European Journal of Operational Research*, Vol. 140 No. 2, pp. 354-372.
- R. N. Shepard, 1982. "Structural Representations of Musical Pitch," in D. Deutsch (ed.): *The Psychology of Music*, Orlando: Academic Press, pp. 343-390.
- M. J. Steedman, 1996. "The Blues and the Abstract Truth: Music and Mentals Models," in (eds.) A. Garnham, J. Oakhill, *Mentals Models in Cognitive Science*. Psychology Press, Hove, pp. 305-318.
- R. Swanson, E. Chew, and A. S. Gordon, 2008. "Supporting Musical Creativity with Unsupervised Syntactic Parsing," in *Proceedings of the AAAI Spring Symposium on Creative Intelligent Systems*, Stanford University, California.
- C. Truchet and P. Codognet, 2004. "Solving Musical Constraints with Adaptive Search," *Soft Computing*, Vol. 8, No. 9, pp. 633-640.
- G. Widmer, 2001. "Using AI and machine learning to study expressive music performance: project survey and first report," *AI Communications*, Vol. 14, No. 3, pp. 149-162.

---

**Elaine Chew** (echew@usc.edu) is the Edward, Frances, and Shirley B. Daniels Fellow at the Radcliffe Institute for Advanced Study at Harvard University in Cambridge, Massachusetts. During 2007-2008, she is on sabbatical from the University of Southern California Viterbi School of Engineering in Los Angeles, California, where she is associate professor of Industrial and Systems Engineering, and of Electrical Engineering.