Preprint from *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning,* Shlomo Argamon, Kevin Burns, and Shlomo Dubnov (Eds.), Berlin: Springer-Verlag, 2010, pp. 45-58.

Style in Music

Roger B. Dannenberg

Because music is not objectively descriptive or representational, the subjective qualities of music seem to be most important. Style is one of the most salient qualities of music, and in fact most descriptions of music refer to some aspect of musical style. Style in music can refer to historical periods, composers, performers, sonic texture, emotion, and genre. In recent years, many aspects of music style have been studied from the standpoint of automation: How can musical style be recognized and synthesized? An introduction to musical style describes ways in which style is characterized by composers and music theorists. Examples are then given where musical style is the focal point for computer models of music analysis and music generation.

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213 USA
rbd@cs.cmu.edu
http://www.cs.cmu.edu/~rbd/

Computers are important in many aspects of music composition, production, distribution, and analysis. In contrast to domains such as natural language, speech, and even images, music rarely has a well-defined meaning, referent, or objective. Consider the sentence, “Tie your shoes.” There is a basic, objective meaning that forms a well-understood command (at least to English-speaking humans). With spoken text and in different contexts, one can imagine all sorts of nuanced versions expressing anger, sympathy, embarrassment, impatience, authority, gentleness, and so on.

Now consider a short melody without words. There is no obvious objective meaning, story, or referent associated with a pure melody. Everything that we enjoy (or not) about the melody has to do with expectations, sound quality, performance nuance, and musical texture. Essentially every aspect of the melody that communicates something to the listener is an aspect of *style.*

In that sense, style is everything in music. So music is a wonderful domain to think about style, but at the same time, style is so broad and vague that we will only make progress if we deconstruct style into more specific concepts. As one might hope, music theory and the many dimensions of music orchestration and performance offer many opportunities to investigate style. In addition to theoretical writings, there are many stylistic concepts that have been modeled with computers and studied more objectively.

In the next section, we will discuss the nature of style in music and talk about how one might describe or characterize musical style. This section is written for the musical novice. If you “don’t know anything about music but you know what you like,” perhaps this section will offer some terminology and help to understand how musicians think about music structure and organization as it relates to style. Section 2 presents a number of computer models of style for both analysis and generation. This section will assume some general knowledge of computer science including data structures, and algorithms.

What Is Musical Style?

In general, "style" means a distinctive quality, form, or type. A more specific definition that certainly applies to music is "a particular manner or technique by which something is done, created, or performed." (Merriam-Webster 2007) In music, the term "style" is used in many ways:

- Historical periods of music are associated with styles. For example, we might say a composition by Mozart is in the Classical style, and one by Bach is in the Baroque style. These styles can be more or less specific: in the recording industry, the term "Classical" is so broad that Mozart and Bach are both "Classical" composers, but a music scholar might speak of "late Classical" or "Neapolitan Baroque."
- Styles are associated with composers. We can speak of composing in the style of Beethoven. In this sense, "style" means "a set of characteristics generally found in the works of a particular composer."
- Performers, especially improvising performers, also have styles. The "ballad style of Miles Davis" refers to characteristics of Miles Davis's performances of ballads. Of course, great classical music players interpret music they perform even if the music is not improvised. One can speak of the expressive style of Itzhak Perlman, for example.
- Style can refer to aspects of musical texture. "Texture" is one of those words like "style" that is very difficult to pin down, and dictionaries do not consider the depth of meaning that texture has for composers. Basically, musical texture is a composite of many aspects of music that one would hear within a second or so. On longer time scales, melodic, rhythmic, and harmonic progressions stand out, but at shorter time scales, we hear timbre (violins?, electric guitars?, saxophones?), very short repeated patterns, many or few different pitches, loudness, and brightness, all of which give a subjective impression we call texture. While composers usually consider "texture" to be something different from style, texture is at least a closely related concept. "Texture" usually refers to sound and the activity of making sound, while "style" is most often used to describe the general impression or intention provided by a texture. We speak of a "tonal style," a "heavy style," or a "big band style," all of which refer to texture-induced impressions. In these examples, the style is not so much the melody, rhythm, or harmony, but the *sound color* in which these elements are embedded.
- Music is often described in emotional terms: exciting, soothing, calm, scary, etc. Sometimes music causes listeners to experience emotions, and other times the listener may recognize an emotion without necessarily experiencing it. Either way, emotional associations are yet another way to describe the style of music.
- Style is often used to mean "genre," yet another difficult-to-define term. A genre is a category of music characterized by a particular style, but a genre can also be influenced by social conventions, marketing, association with a particular artist, and other external influences. Still, it is common to refer to something as "rock style," or "bebop style."

All of these definitions are related to the underlying idea that there are important characteristics of music that we perceive as common or related across certain collections – the work of a composer, the output of some historical period, or music of some genre. Musicians study the elements of music in detail and are familiar with ways in which these elements can be varied, giving rise to different styles. It is interesting that non-musicians can also perceive styles with great sensitivity, often with no ability to describe characteristics or differences. For these listeners, I will offer some terminology and discussion through examples. Do not expect to become a musical expert, and do not believe that experts have a complete formal model of style, but hopefully this discussion will explain some of the ways musical style can be treated objectively. For more details on music terminology, concepts, and history, the *New Grove Dictionary of Music and Musicians* (Sadie and Tyrrell, 2001) is an excellent reference.

An Example: Baroque vs. Classical Style

A good way to learn about musical style is to examine the differences between two well-known styles. In this section, we will compare *Baroque* and *Classical* styles. The Baroque period extends from about 1600 to 1750 and includes music by Monteverdi, Bach, Handel, and Vivaldi.

As noted above, "classical" is sometimes used to refer to a broad range of styles sometimes referred to as "Western art music," which includes Renaissance, Baroque, Classical, Romantic, and many modern styles. But to experts, Classical music is music from the Classical period, approximately 1750 to 1800. The most celebrated Classical composers are Haydn, Mozart, and Beethoven in his early years.

Baroque and Classical music differ along many dimensions as listed in Table 1. It should be noted that there are few absolutes in music, and certainly there are exceptions to every rule. However, scholars generally agree that the characteristics described here are important, real differences between the two styles. These differences are now presented in more detail.

Table 1. Characteristics of Baroque and Classical Styles

Baroque	Classical
Contrapuntal	Homophonic
Ornamented	Internal Structure
Frequent, less meaningful modulation	Modulation becomes structural element
Single vivid feeling	Range of emotions
Constant intensity throughout	Dramatic climax and resolution

The first difference is contrapuntal vs. homophonic writing for multiple voices. Here, "voice" is used in a technical sense that means a human voice *or* an instrument, so the difference is how composers combine multiple simultaneous sounds. Homophonic writing emphasizes harmony and a single dominant melody. Typically, the highest voice carries the melody, and concurrently with each melody note, the other voices sing harmonizing pitches. Most church hymns are homophonic. In contrast, contrapuntal writing emphasizes melodic motion over harmony. It is as if each voice is singing its own melody. Often one voice will sing a short melodic sequence and then hold a steady pitch while another voice sings an "answer" in the form of another melodic sequence. Two, three, or more melodies are thus intertwined to create what is called counterpoint. In more objective terms, at least one feature of contrapuntal writing is that fewer notes begin synchronously compared to homophony.

In Baroque writing, there is an emphasis on ornamentation, typically short, fast notes inserted above and below at the beginning of a "normal" note in the melody. The trill, where the pitch alternates between the melody note and the next one above, is another ornament. Some ornaments and their interpretations are illustrated in Figure 1. If Baroque style tends to "decorate" melodic lines, the Classical style is plainer, with an emphasis on developing ideas and formal structure. For example, the sonata form that appears in the classical period is based on a structure consisting of two themes, their development, a return to the themes, and a conclusion.

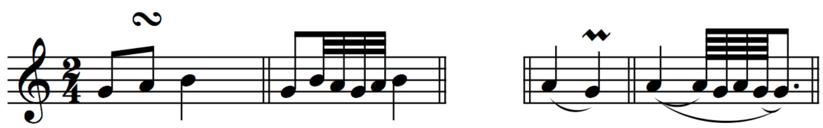

Figure 1. Some musical ornaments and their interpretations: Turn (left) and Mordent (right).

Baroque and Classical music is based on musical scales. For example, the white keys of the piano (omitting the black ones) are used in the scale of C-Major. Any music can be *transposed* by adding an offset to each piano key. For example, if every note is played 7 keys to the right on the keyboard (counting white notes and black notes), the result will sound very similar, but the music has been translated to a new location. This is called a modulation. Modulation in Baroque music occurs frequently, but the modulations do not have much significance. In Classical music, modulations are often carefully resolved: a modulation up 7 steps is likely to come back down 7 steps, setting up an expectation in the listener. In the sonata form, modulation is used to announce the introduction of the second theme, and modulation is often carefully coordinated with other musical structure.

In terms of feeling, Baroque music typically portrays a single emotion at least through an entire movement or major section of a composition. Classical music is more likely to progress from one feeling to another in a narrative style, exploring a range of emotions. While Baroque music is more likely to maintain a feeling with a steady intensity, Classical music often develops into a climax of tension, excitement, and just plain loudness, and then settles into a state of calm and resolution. We will consider later how something as abstract as "feeling" in music can be quantified and studied objectively.

To summarize, Baroque and Classical music differ along a number of dimensions. These differences can be difficult to formalize and even to describe to non-musicians, but at least it should be clear that most music lovers, with a little experience, can at least recognize these styles. Music scholars can go further by describing differences in specific terms. Music is one area where "style" has been deeply studied and where there are many examples and analyses of different styles.

Style in Popular Music

What makes a modern popular musical style? There are so many emerging styles that most people are not even familiar with many of them. What distinguishes Black Metal, Death Metal, Doom Metal, Hair Metal, and Power Metal? (Hint for Classical purists: these are rock styles.) And if you are a hard-core metal enthusiast for whom these terms are familiar, what distinguishes Be-Bop, Hard-Bop, and Post-Bop? (Hint: think jazz.) Rather than tackle these questions specifically, we will look at some general characteristics of music. Style, especially in popular music, includes an important sociological component, so we should not expect style to be purely a matter of how something sounds. The composer, performer, geographic region, marketing, and public perception have an important influence on how music is categorized.

As always, popular music style has many meanings and interpretations. We could talk about singing style, genre, rhythmic feel, dance styles, and others. Without getting too specific, imagine scanning a radio dial looking for a favorite style of music. Experiments by Perrot and Gjerdigen (1999) indicate that we can recognize style in a fraction of a second. What elements of music allow us to quickly determine if we have found something in the style we are looking for?

In popular music, one very important element is the instrumentation. If we hear nothing but guitar, bass, and drums, this might be hard rock, but if we hear a saxophone and trumpet, this might be a blues band. Strings might indicate pop music or a rock ballad. In addition to instruments, the quality of instrumental sounds is important. An acoustic guitar or pure guitar sound might indicate soft rock or country music, while a highly distorted electric guitar is more typical of heavy metal. As we shall see, automatic genre classifiers can be based purely on the average frequency spectrum of music audio.

Vocal quality and the number of vocalists (with harmony) also tell us something. Steady clear vocals, spoken words (as in Rap), screaming, falsetto singing, and the use of pitch inflections and vibrato could all be described as vocal styles, and all tell us something about the style of music. You would not hear an operatic voice singing country music or four-part harmony singing on a techno track.

Rhythm is very important because most popular music is very rhythmic. There are rhythmic patterns associated with different styles of music as well as with dance styles. Rock is characterized by strong beats in groups of four, with accents on 2 and 4: one-TWO-three-FOUR. Compare this to Reggae, which also follows the general rock pattern, but often with a slower tempo, a subdivision of beats (e.g. one-and-TWO-and-three-and-FOUR-and), and emphasis on the rhythmic bass patterns. Reggae is a good example of the importance of not only the rhythm but *how* the rhythm is established by different instruments including different types of drums and other percussion.

Computational Approaches to Music Style

In recent years, many advances have been made in the analysis of musical style and the generation of music according to certain styles. Some of these advances can be attributed to advances in statistical machine learning, which seems to be well-suited to the capture of style information where data is more representative of trends than hard-and-fast rules.

Learning to Recognize Improvisational Styles

A standard task is a forced-classification of music into one of a set of style categories. The general approach is seen in Figure 2. The input is music data in the form of audio. The first step is to extract features from the audio signal. While it is theoretically possible that a system could learn to classify styles directly from digital audio signals, this is not practical. Instead, we perform some analysis on the sound to obtain a small set of abstract features that will hopefully contain useful information for discriminating styles. Next, a classifier is used to estimate the style of the sample. The classifier can be based on any number of machine learning models. For this discussion, we will only be concerned with the general nature of these systems. Basically, a classifier begins with a number of labeled examples called the *training set.* Each example contains a set of features obtained from an excerpt of music and a *label*, which gives the correct style for this excerpt. There may be many thousands of examples. From the examples, the classifier learns to output the correct label given a set of feature values. Learning is usually accomplished by iteratively adjusting parameters within the classifier to improve its performance on the training set. For details, consult a textbook on machine learning (Mitchell, 1997).

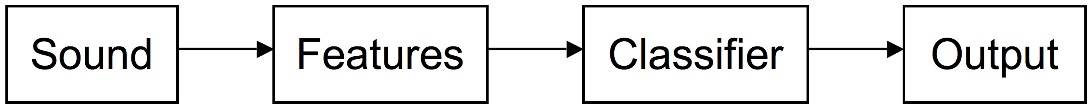

Figure 2. Style Classification

The features obtained from the music are critical to machine learning, especially when something as abstract as style must be determined from something as concrete as an audio waveform. Surprisingly, style classification does not always require high-level features like pitch and rhythm, but let us start there anyway.

Belinda Thom and I created what may be the first music style classification system (Dannenberg, Thom, and Watson 1997). Our goal was to detect different improvisational styles. There is no list of standard improvisation styles, so at first glance, this problem may seem to be poorly defined. How can we know if the classifier works? We avoided the need for absolute, objective definitions by letting the improviser define style categories and give training examples. A classifier was then trained to recognize these categories. To test the classifier, a computer would display the name of a style category, the improviser would play in that style, and the classifier would attempt to recognize the style. If the classifier output matched the displayed style category, then we claimed that recognition had indeed occurred. The beauty of this experiment is that, to the improviser, the different styles remained purely subjective concepts, and yet we were able to make objective ratings of classifier performance.

For this classifier, the input was assumed to be monophonic audio, that is, the sound of one instrument. This sound was analyzed to extract pitches, note onset times, note durations, overall note amplitude, and time-varying amplitude and pitch. We organized this data into 5-second “windows” because we wanted a system that could respond to changes in style. There is a tradeoff here: longer windows give more reliable statistical information but reduce the responsiveness to changes. We guessed that five seconds would be long enough to collect useful information and short enough to be useful in real-time interactive music systems that would respond to performers’ styles.

The features we used included the average (mean) pitch, number of notes, average loudness, average duration, and average duty cycle. The “duty cycle” is the fraction of total time that a note is sounding. In addition, we computed standard deviations of pitch, duration, and loudness. Additional statistics were computed on pitch changes and loudness changes. All of these statistics refer to a 5-second period of playing. We analyzed a 5-second period starting at every second, so the windows overlapped both in training and in testing the classifier.

We worked with a set of 4 distinctive styles that were labeled *lyrical, pointillistic, syncopated,* and *frantic.* These categories are easily distinguishable by humans and our machine classifiers, which in one case recognized over 99% of the test cases correctly. It is interesting that *syncopated*, which literally means placing notes on up-beats (between beats) rather than down-beats, can be recognized without the need to detect tempo and beats. Evidently, there is a characteristic manner of playing syncopations that manifests itself through other features, including the ones that we detected.

We also tested the system on a more difficult set of 8 styles. In addition the four listed above, we used *blues, quote, high,* and *low.* In practice, these "styles" are not mutually exclusive. One could easily play "lyrical blues" or "high and frantic." Also, *quote* means to play a familiar tune from memory. It would seem impossible that the computer could detect quotes without a memory of popular tunes, but in fact the computer recognition of this category was still better than chance. As with *syncopated*, we believe that quotes are played in a manner that can be detected through low-level features. Overall, the computer recognized 90% of these styles correctly, whereas random guessing would get 12.5% correct.

Questions these experiments cannot answer include "does *frantic* have a 'true' nature?" and "what is it?" Perhaps like *syncopation*, there is a deeper human notion of each of these styles, and we are just measuring superficial features that happen to be correlated with the "true" styles. Our simple features certainly do not capture all the kinds of musical information that human listeners hear and process, so it does seem likely that *frantic* and other categories *do* have deeper and richer meanings and associations than indicated by our computer models. However, it is still interesting how well simple machine models can learn to recognize stylistic differences in musical performance.

Genre Classification

The same basic ideas of feature extraction, labeled training examples, and machine learning have been used for the task of automatic genre classification. The goal here is to label music audio according to *genre*, such as rock, pop, country, jazz, and classical. Much of the research in this area investigates features.

Surprisingly (at least to this author), the average spectrum is quite effective for genre classification. How could this be? After all, the average spectrum does not convey any information about beats, harmonic progressions, or melodic shape. Aren't these the things that differ among different genres? It turns out that, for example, rock music has a very broad spectrum with a lot of energy at high frequencies (think penetrating, distorted guitars and snare drums) compared to classical music. Thus, a binary rock/classical music detector should be quite simple to create.

Other spectrally related correlations can be found. For example, rock music tends to stay in one key, whereas jazz and classical music are more likely to modulate to different keys. This affects the distribution of energy in the frequency spectrum. Trying to apply reasoning and rules to genre classification might be partially successful, but a much better approach is to use machine learning, where hundreds or thousands of features can be considered systematically. Typically, researchers use more than just the spectrum, although for the most part, the features do not involve any kind of high-level music perception. One might think that by detecting high-level features such as tempo, meter, orchestration, and even music transcription (what pitches are played when), that genre classification could be improved. In practice, these higher-level features are very difficult to detect, and detection errors limit their usefulness in genre classification.

Some features that can be obtained automatically include average amplitude, beat strength, and harmonicity. Average amplitude is just a measure of the audio energy within each short time window, say, 10ms. A histogram of amplitude values or just the standard deviation can tell whether the music stays mainly at the same level, typical of popular music, or has a wide dynamic range, typical of classical music. The beat strength is a measure of whether amplitude variations tend to be periodic or not. A measure of tempo and tempo variation might be even more useful, but tracking tempo reliably is a difficult problem. Harmonicity is an indication of how well the spectrum can be modeled as a set of harmonically related partials. Harmonic spectra arise from single tones or sets of harmonically related tones. Harmonicity may also be viewed as periodicity in the frequency domain.

All of these features help to capture audio qualities that give clues about genre. “Genre” is not a very specific term, so “genre classifiers” are also useful for other tasks. Depending on the training set, a genre classifier can be trained to recognize genre (Tzanetakis and Cook 2002), recognize different styles of dance music (Dixon et al. 2003), find music “similar” to an example (Berenzweig et al. 2004), or find music that I like as opposed to music I do not like. Classifiers are potentially useful for music search, but there is considerable interest in using classifiers and similarity metrics to create music recommendation services (Magno and Sable 2008), generate playlists (Pauws et al. 2006), and organize personal music libraries (Torrens et al. 2004).

Markov Models

Many more techniques are available to model music and style when the data consists of discrete (or symbolic) notes as opposed to audio recordings. Discrete representations of music are generally in the form of a list of notes, where a note is a tuple consisting of pitch, starting time, duration, and loudness. Pitch is often represented using an integer representing the corresponding key on a piano, while time and duration can be represented either in terms of seconds or in beats. Other attributes are also possible. The Standard MIDI File format is a common example of a discrete music representation. Standard MIDI Files can be used to control music synthesizers, thus a MIDI file can be “played” on most computers (Rothstein 1995).

An interesting way to think about style is to look for commonalities in discrete music data. One model, proposed by Alamkan, Birmingham, and Simoni (1999), is the Markov Model applied to what these researchers call “concurrencies.” A concurrency is a set of pitches and a duration. To extract concurrencies from music, the music is first segmented at every time point where a note begins or ends. (See Figure 3.) It follows that within a segment, some set of pitches is sounding, and the set does not change within the segment. (Otherwise, the segment would be divided further where the note starts or stops.) The space of all concurrencies is huge: it is $2^{|p|}\times|d|$, where $|p|$ is the number of possible pitches, and $|d|$ is the number of possible durations. In practice, a given piece of music will contain relatively few concurrencies, and due to repetition and similarity within a given composition, many concurrencies will occur often.

Alamkan, Birmingham, and Simoni used Markov Models of concurrencies to model Beethoven and Mozart piano sonatas. A Markov Model is a set of states (in this case, each concurrency is a state) and transition probabilities. Transitions probabilities are estimated based on the number of actual transitions observed in a piano sonata or set of sonatas. A measure of similarity between Markov Models was developed. Using this approach, it can be shown that early Beethoven sonatas are more similar to Mozart sonatas than to late Beethoven sonatas. This is consistent with the music history view that Beethoven’s early music is Classical (like Mozart) and his later works belong to the period known as “Romantic.”

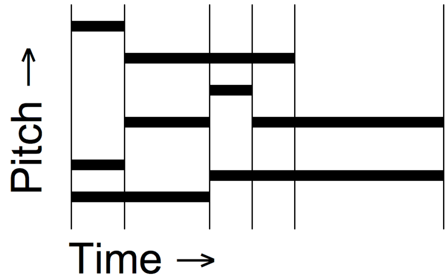

Concurrency #: 1 2 3 4 5

Figure 3. "Piano-roll" depiction of music with concurrencies marked.

Markov Models can also be used for music *generation.* One approach is to simply follow transitions according to their probabilities, assembling the resulting sequence of concurrencies into a new composition. The resulting music is often strikingly good at the local level, because every transition is not only plausible but was used at least once in a master work. On the other hand, Markov Models have no higher-level sense of direction or form, so the music typically meanders without any sense of purpose.

Cope's Experiments in Musical Intelligence

One of the most successful attempts to capture and model Western musical styles is the work of David Cope referred to as EMI, or Experiments in Musical Intelligence (Cope, 1991). A great deal has been written about this work by Cope and others, so we will give only a brief summary of the nature of this work.

One interesting idea in EMI is the use of short patterns of melody to represent the essence of what characterizes composers and compositions. Within a piece, recurring patterns tell us about themes and motives that are specific to the composition. When many compositions by the same composer are analyzed, we can discover patterns that are characteristic of the composer, especially if we first rule out the patterns that are specific to individual pieces. These "composer patterns" can be used as building blocks for a new piece in the style of the composer.

Another important idea in EMI is the use of Augmented Transition Networks, or ATNs, a type of formal grammar. ATNs were originally developed to describe natural language, including not only grammatical syntax, but semantic constraints as well. In EMI, ATNs are used to guide the generation of music so that there is some global coherence, not just local plausibility as produced by Markov Models. Overall, EMI works by reassembling patterns from example music to form new compositions that are stylistically similar.

Evaluation of this sort of work is very difficult. There are no standards to measure whether a piece is really in the style of some composer. For the most part, this work must be evaluated by subjective listening. Cope has set up a sort of "musical Turing test" where listeners are asked to rate a phrase as original Mozart or "artificial" Mozart. Many people are unable to distinguish the two, but experts seem to have little difficulty. A critic might argue that listeners are confused by the fact that at least fragments of what they hear really *were* composed by Mozart and only recombined to form the "artificial" Mozart examples. On the other hand, EMI shows that Mozart himself reuses small "composer patterns," so any convincing Mozartean music arguably *must* use these patterns.

The debate over computers, creativity, intelligence, and style is certainly not over. Cope makes the interesting point that "artificial" music composed by EMI sounds much more convincing when performed by humans. The next section explores some stylistic aspects of musical performance.

Emotion and Expression in Music

The previous sections on Markov Models and EMI examine style as an aspect of music composition. In those studies, we looked for style in the selection of pitches, rhythms, and their transitions and organization. At least with Western music, which is strongly tied to music notation, the notated "composition" is viewed as separate from the *interpretation* of the composition by a performer. Performance interpretations also have style. A performance is not a precise, mechanical rendition of the printed notes, which would sound very unmusical. A performer adds slight variations to tempo, loudness, note "shape" or articulation, slight pauses, vibrato, and many other details that are at best sketchy in the printed music. Through all these details, a performer brings life to the printed page.

Researchers have discovered that performance nuance can impart emotional meaning to music. Although we often think of a particular song as being happy or sad, it turns out that the performance has as much to do with the emotional impact as the composition. Given a "neutral" melody, one can perform the melody in such a way as to express joy, sadness, anger, or serenity, for example. Juslin and Sloboda (2001) have published an excellent reference on emotion in music.

One of the more interesting findings is that computers can alter music performances to express different emotions. A good example is the work of Bresin and Friberg (2000). This work began with Sundberg's pioneering work on expressive music performance (1988). Sundberg studied how to make "musical" performances from music in standard notation. The input is assumed to be a discrete representation where note times and durations are quantized to exact beat boundaries, no tempo variation is indicated, and no loudness variation or note articulations are specified. In other words, the input contains the information that a human performer would encounter in ordinary printed music. The goal is to add variation in loudness and timing to make the performance more musical. To accomplish this, rules were developed. For example, one rule says that if the melody leaps upwards, add a small pause before the leap. Another rule says that if the melody goes up by a run of small steps, the notes should get progressively faster and louder.

Getting back to emotion, at least one way to impart emotion to a performance is to change some overall performance parameters and also to change the weights on these performance rules. For example, to make music sound angry, increase the overall loudness, decrease note duration so that there is more space between notes, and decrease the effect of rules that vary tempo. To impart the feeling of fear, increase the effect of rules that vary tempo, and play the music slower and softer overall. There are many more rules, so these prescriptions only touch on a much larger set of carefully worked out rules and parameters.

To evaluate Bresin's results, listeners were asked to choose which of 6 emotions were strongest in different computer renditions of well-known folk songs. The renditions were generated automatically by Bresin's system. The human subjects were able to identify the "correct" emotions significantly above chance levels. With only simple and short folk melodies, it is not clear whether subjects truly felt these emotions or whether the "emotional intent" was recognized but not felt. Of course, the most interesting goal is to give computers the ability to create "real" emotion. There are plenty of opportunities for emotionally expressive music, especially in interactive settings such as video games. Even cell phone ring tones may some day be used to express the caller's emotion. It seems likely that future computer music systems will have the ability to give emotional performances.

Summary and Conclusion

As in other disciplines, *style* in music is used in many ways with many shades of meaning. A musical style can refer to a period in history, a composer, a performer, a texture, a genre, or an emotion. While style tends to include all those ill-defined characteristics that we cannot quite grasp, musical styles have been carefully studied, partly because style and subjective impressions are at the heart of musical communication. We saw in the comparison of Baroque and Classical styles that the differences can be described quite explicitly.

Style can be recognized and even learned by computer programs. We considered style and genre classifiers that use standard machine learning techniques to discriminate different styles on the basis of a wide range of features that can be computed from audio signals. Discrete representations have been used, for example with Markov models, to represent and compare musical styles. Markov models have shown some promise for stylistic music generation. Music patterns can be detected and recombined to "compose" music in the style of a particular composer, as illustrated by David Cope's EMI. Music can also be performed with consideration for style. For example, programs can perform melodies in ways that convey a wide range of emotions that are understood by human listeners.

Music analysis by computer is a very active research field, in part because of the enormous amounts of musical data being stored in databases and either sold or shared over the Internet. This has motivated research that will allow computers to understand and compare music based on music *content* as opposed to titles, artist names, and other textual labels. In the future, we expect to see more computer applications that make use of music style for automatic music analysis, music search, and even music generation.

References

Alamkan C, Birmingham (1999) Stylistic structures: An initial investigation of the stochastic genration of tonal music. University of Michigan Technical Report CSE-TR-395-99

Berenzweig A, Logan B, Ellis D, Whitman B (2004) A large-scale evaluation of acoustic and subjective music similarity measures. Computer Music Journal 28(2):63-76

RBresin, R. and, A Friberg (2000) Emotional coloring of computer-controlled music performance. Computer Music Journal. 24(4):44-62

Cope D (1991) Computers and musical style. (The Computer Music and Digital Audio Series, 6.) A-R Editions: Middleton, Wisconsin

Dannenberg RB, Thom B, Watson D (1997) A machine learning approach to musical style recognition. In: 1997 international computer music conference. International Computer Music Association, San Francisco, 344-347

Dixon S, Pampalk E, Widmer G (2003) Classification of dance music by periodicity patterns. In: Fourth international conference on music information retrieval (ISMIR 2003), 159-165

Juslin PN, Sloboda JA (2001) Music and emotion. Oxford University Press, London

Magno T, Sable C (2008) A comparison of signal-based music recommendation to genre labels, collaborative filtering, musicological analysis, human recommendation, and random baseline. In: ISMIR 2008 Proceedings of the 9th international conference on music information retrieval, 161-166

Merriam-Webster (2007) Merriam-Webster Online. http://www.m-w.com/dictionary/style. Accessed 26 November 2008

Mitchell Tom (1997) Machine learning. McGraw Hill, New York

Pauws S, Eggen B (2002) PATS: Realization and user evaluation of an automatic playlist generator. In: Proceedings of the 3rd international conference on music information retrieval (ISMIR'02), 222-230

D and R Gjerdigen (1999) Scanning the dial: an exploration of factors in identification of musical style (abstract only.) In: Proceedings of the society for music perception and cognition, 88

Rothstein J (1995) MIDI: A comprehensive introduction, 2nd edn. A-R Editions, Middleton, Wisconsin

Sadie S and J Tyrell (2001) The new grove dictionary of music and musicians, 2nd edn. Oxford University Press, Oxford

J (1988) Computer synthesis of music performance. In: Sloboda JA (ed) Generative processes in music. The psychology of performance, improvisation, and composition. Clarenden Press, Oxford

Torrens M, Hertzog P, Arcos JL (2004) Visualizing and exploring personal music libraries. In: 5th international conference on music information retrieval, 421-424

Tzanetakis G, Cook P (2002) Musical Genre Classification of Audio Signals. IEEE Transactions on Speech and Audio Processing, 10(5):293-302