Wei Chai



© DIGITALVISION, © ARTVILLE (CAMERAS, TV, AND CASSETTE TAPE) © STOCKBYTE (KEYBOARD)

Semantic Segmentation and Summarization of Music

Methods based on tonality and recurrent structure

istening to music and perceiving its structure are fairly easy tasks for humans, even for listeners without formal musical training. However, building computational models to mimic this process is a complex problem. Furthermore, the amount of music available in digital form has already become unfathomable. How to efficiently store and retrieve digital content has become an important issue. This article presents our research on automatic music segmentation and summarization from audio signals. It will inquire scientifically into the nature of human perception of music and offer a practical solution to difficult problems of machine intelligence for automated multimedia content analysis and information retrieval. Specifically, three problems will be addressed: segmentation based on tonality analysis, segmentation based on recurrent structural analysis, and summarization (or music thumbnailing).

Successful solutions to the above problems can be used for Web browsing, Web searching, and music recommendation. Some previous research has already attempted to solve some similar problems. For segmentation, some research attempted to segment musical signals by detecting the locations where significant changes in statistical properties occur [2], which typically has nothing to do with the high-level structure. There has also been research trying to consider semantic musical structure for segmentation. For example, Sheh [9] proposed using expectation maximization (EM)-based hidden Markov model (HMM) for chord-based segmentation. For summarization, Dannenberg [8] presented a method to automatically detect the repeated patterns of musical signals using self-similarity analysis and clustering. Logan [13] attempted to use a clustering technique or HMM to find key phrases of songs. Bartsch [1] used the similarity matrix proposed by Foote [10], [11] and chromabased features for music thumbnailing. A variation of the similarity matrix was also proposed for music thumbnailing [15]. Previous research typically assumes that the most repeated pattern is the most representative part of music. There has been little research aimed at generating a global recurrent structure of music and a semantic segmentation based on this structure.

CHROMAGRAM REPRESENTATION

The chromagram, also called the pitch class profile (PCP) feature, is a frame-based representation of audio, very similar to short-time Fourier transform (STFT). It combines the frequency components in STFT belonging to the same pitch class (i.e., octave folding) and results in a 12-dimensional representation, corresponding to C, C#, D, D#, E, F, F#, G, G#, A, A#, and B in music, or a generalized version of 24-dimensional representation for higher resolution and better control of noise floor [6].

Specifically, for the 24-dimensional representation, let $X_{\text{STFT}}[K, n]$ denote the magnitude spectrogram of signal $x[n]_{,,}$ where $0 \le K \le \text{NFFT} - 1$ is the frequency index and NFFT is the FFT length. The chromagram of x[n] is

$$X_{\text{PCP}}[\tilde{K}, n] = \sum_{K: \ P(K) = \tilde{K}} X_{\text{STFT}}[K, n].$$
(1)

The spectral warping between frequency index K in STFT and frequency index \tilde{K} in PCP is

$$P(K) = [24 \cdot \log_2(K/\text{NFFT} \cdot f_s/f_1)] \mod 24, \tag{2}$$

where f_s is the sampling rate and f_1 is the reference frequency corresponding to a note in the standard tuning system, for example, musical instrument digital interface (MIDI) note C3 (32.7 Hz). For the following two segmentation tasks, chromagram will be employed as the representation.

MUSIC SEGMENTATION BASED ON TONALITY ANALYSIS

This section describes an algorithm for detecting the key (or keys) of a musical piece. Specifically, given a musical piece (or a part of it), the system will segment it into sections based on key change and identify the key of each section. Note that here we want to segment the piece and identify the key of each segment at the same time. A simpler task could be, given a segment of a particular key, to detect the key of it.

In the following, the task of key detection will be divided into two steps. 1) Detect the key without considering its mode. (For example, both C major and A minor will be denoted as key 1, C# major and A# minor will be denoted as key 2, and so on. Thus, there could be 12 different keys in this step.) 2) Detect the mode (major or minor).

The task is divided in this way because diatonic scales are assumed and relative modes share the same diatonic scale. Step 1 attempts to determine the height of the diatonic scale. And again, both steps involve segmentation based on key (mode) change as well as identification of keys (modes).

The model used for key change detection should be able to capture the dynamic of sequences and incorporate prior musical knowledge easily since a large volume of training data is normally not available. We propose to use HMMs for this task because HMM is a generative model for labeling structured sequences and satisfying both of the above properties. The hidden states correspond to different keys (or modes). The observations correspond to each frame represented as 24-dimensional chromagram vectors. The task will be decoding the underlying sequence of hidden states (keys or modes) from the observation sequence using the Viterbi approach [16].

The parameters of HMM that need to be configured include:

- The number of states *N* corresponding to the number of different keys (=12) or the number of different modes (=2), respectively, in the two steps.
- The state transition probability distribution $A = \{a_{ij}\}$ corresponding to the probability of changing from key (mode) *i* to key (mode) *j*. (Thus, A is a 12×12 matrix in step 1 and a 2×2 matrix in step 2.)
- The initial state distribution $\Pi = {\pi_i}$ corresponding to the probability at which a piece of music starts from key (mode) *i*.
- The observation probability distribution $B = \{b_j(v)\}$ corresponding to the probability at which a chromagram v is generated by key (mode) j.

Due to the small amount of labeled audio data and the clear musical interpretation of the parameters, we will directly incorporate the prior musical knowledge by empirically setting Π and A as follows:

$$\Pi = \frac{1}{12} \cdot 1,\tag{3}$$

where **1** is a 12-dimensional vector in step 1 and a two-dimensional vector in step 2. This configuration denotes equal probabilities of starting from different keys (modes).

$$\mathbf{A} = \begin{bmatrix} \text{stayprob} & b & \dots & b \\ b & \text{stayprob} & \dots & b \\ b & b & \dots & b \\ b & b & \dots & \text{stayprob} \end{bmatrix}_{d \times d}, \quad (4)$$

where *d* is 12 in step 1 and 2 in step 2. stayprob is the probability of staying in the same state and stayprob + $(d - 1) \cdot b = 1$. For step 1, this configuration denotes equal probabilities of changing from one key to a different key. It can be easily shown that when stayprob gets smaller, the state sequence becomes less stable (changes more often). In our experiment, stayprob will be varying within a range (e.g., [0.9900 0.9995]) in step 1 to see how it impacts the performance; it will be empirically set to $1 - 10^{-20}$ in step 2.

For observation probability distribution, Gaussian probabilistic models are commonly used for modeling observations of continuous random vectors in HMM. Here, however, the cosine distances between the observation (the 24-dimensional chromagram vector) and predefined template vectors were used to represent how likely it was that the observation was emitted by the corresponding keys or modes, i.e.,

$$b_j(v) = \frac{v.\theta_j}{||v||.||\theta_j||},\tag{5}$$

where θ_j is the template of state *j* (corresponding to the *j*th key or mode). The advantage of using cosine distance instead of Gaussian distribution is that the key (or mode) is more correlated with the relative amplitudes of different frequency components rather than the absolute values of the amplitudes.

$$\theta_j = r(\theta_1, 2 \cdot (j-1)), \tag{6}$$

$$\beta = r(\alpha, k), \ s.t. \ \beta[i] = \alpha[(k+i) \mod 24], \tag{7}$$

where j = 1, 2, ..., 12 and i, k = 1, 2, ..., 24. Let us also define 24 mod 24 = 24.

MUSIC SEGMENTATION BASED ON RECURRENT STRUCTURAL ANALYSIS

Music typically has a recurrent structure. This section describes research into automatic identification of the recurrent structure

of music from acoustic signals. Specifically, an algorithm will be presented to output structural information, including both the form (e.g., AABABA) and the boundaries indicating the beginning and the end of each section. It is assumed that no prior knowledge about musical forms or the length of each section is provided and that the restatement of a section may have variations (e.g., different lyrics, tempos). These assumptions require both robustness and efficiency of the algorithm.

REPRESENTATION FOR SELF-SIMILARITY ANALYSIS

For visualizing and analyzing the recurrent structure of music, Foote [10], [11] proposed a representation called self-similarity matrix. Each cell in the matrix denotes the similarity between a pair of frames in the musical signal. Here, instead of using similarity, we will use distance between a pair of frames, which results in a distance matrix (DM). Specifically, let $V = v_1v_2, \ldots, v_n$ denote the feature vector sequence of the original musical signal *x*. It means we segment *x* into overlapped frames x_i and compute the feature vector v_i of each frame (e.g., chromagram). We then compute the distance between each pair of feature vectors according to some distance metric and obtain the DM. Thus,

$$DM(V) = [d_{ij}] = [||v_i - v_j||],$$
(8)

where $||v_i - v_i||$ denotes the distance between v_i and v_i .

Since distance is typically symmetric, i.e., $||v_i - v_j|| = ||v_j - v_i||$, the DM is also symmetric. One widely used definition of distance between vectors is based on cosine distance:

$$||v_i - v_j|| = 0.5 - 0.5 \cdot \frac{v_i \bullet v_j}{|v_i||v_j|},\tag{9}$$

where we normalized the original definition of cosine distance to range from 0 to 1 instead of 1 to -1 to be consistent with the nonnegative property of distance. If we plot the DM, we can often see the diagonal lines in the plot, which typically correspond to repetitions. Some previous research attempted to detect these diagonal patterns for identifying repetitions. However, not all repetitions can be easily seen from this plot due to variations of the restatements.

DYNAMIC TIME WARPING FOR MUSIC MATCHING

The above section showed that when part of the musical signal repeats itself nearly perfectly, diagonal lines appear in the DM or its variation representations. However, if the repetitions have numerous variations (e.g., tempo change, different lyrics), which is very common in all kinds of music, the diagonal patterns will not be obvious. One solution is to consider approximate matching based on the self-similarity representation to allow flexibility of repetitions, especially tempo flexibility. Dynamic time warping was widely used in speech recognition for similar purposes. Previous research has shown that it is also effective for music pattern matching [18]. Note that dynamic time warping is often mentioned in the context of speech recognition, where a technique similar to dynamic programming is cited for approximate string matching, and R the distance between two strings based on this technique is T often called edit distance.

Assume we have two sequences, and we need to find the match between the two sequences. Typically, one sequence is the input pattern $(U = u_1u_2, \ldots, u_m)$ and the other $(V = v_1v_2, \ldots, v_n)$ is the one in which to search for the input pattern. Here, we allow multiple appearances of pattern U in V. Dynamic time warping utilizes the dynamic programming approach to fill in an $m \times n$ matrix WM based on (10). The initial condition (i = 0 or j = 0) is set based on Figure 1

$$DM[i, j] = \min \begin{cases} DM[i-1, j] + c_D[i, j], & (i \ge 1, j \ge 0) \\ DM[i, j-1] + c_I[i, j], & (i \ge 0, j \ge 1) \\ DM[i-1, j-1] + c_S[i, j], (i, j \ge 1) \end{cases}$$
(10)

where c_D is the cost of deletion, c_I is the cost of insertion, and c_S is the cost of substitution. The definitions of these parameters are determined differently for different applications. For example, we can define $c_S[i, j] = ||u_i - v_j||$ and $c_D[i, j] = c_I[i, j] = 1.2 \cdot c_S[i, j]$ to penalize insertion and deletion based on the distance between u_i and v_j . We can also define c_D and c_I to be some constant.

The last row of matrix WM (highlighted in Figure 1) is defined as a matching function r[i] (i = 1, 2, ..., n). If there are multiple appearances of pattern U in V, local minima corresponding to these locations will occur in r[i]. We can also define the overall cost of matching U and V (i.e., edit distance) to be the minimum of r[i], i.e., $||U - V||_{DTW} = \min_i \{r[i]\}$. In addition, to find the locations in V that match pattern U we need a traceback step. The trace-back result is denoted as a trace-back function t[i] recording the index of the matching point. The time complexity of dynamic time warping is O(nm), corresponding to the computation needed for filling up matrix WM.

RECURRENT STRUCTURAL ANALYSIS

Assuming that we have computed the feature vector sequence and the DM, the algorithm follows four steps, which will be explained in the following four sections. All the parameter configurations are tuned based on the experimental corpus that will be described in the "Experiment and Evaluation" section.

PATTERN MATCHING

In the first step, we segment the feature vector sequence (i.e., $V = v_1v_2 \dots v_n$) into overlapped segments of fixed length l (i.e., $S = S_1S_2 \dots S_m$; $S_i = v_{k_i}v_{k_i+1} \dots v_{k_i+l-1}$; e.g., 200 consecutive vectors with 150 vectors overlap) and compute the repetitive property of each segment S_i by matching S_i against the feature vector sequence starting from S_i (i.e., $V_i = v_{k_i}v_{k_i+1} \dots v_n$) using dynamic time warping. We define the cost of substitution c_S to be the distance between each pair of vectors. It can be obtained directly from the DM. We also define the costs of deletion and insertion to be some constant: $c_D[i, j] = c_I[i, j] = a$ (e.g., a = 0.7). For each matching between S_i and V_i , we obtain a matching function $r_i[j]$.

REPETITION DETECTION

This step detects the repetition of each segment S_i . To achieve this, the algorithm detects the local minima in the matching function $r_i[j]$ for each *i*, because typically a repetition of segment S_i will correspond to a local minimum in this function.

There are four predefined parameters in the algorithm of detecting the local minima: the width parameter w, the distance parameter d, the height parameter h, and the shape parameter p. To detect local minima of $r_i[j]$, the algorithm slides the window of width w over $r_i[j]$. Assume the index of the minimum within the window is j_0 with value $r_i[j_0]$, the index of the maximum within the window but left to j_0 is $j_1(i.e., j_1 < j_0)$ with value $r_i[j_1]$, and the index of the maximum within the window but left to j_0 is j_2 (i.e., $j_2 > j_0$) with value $r_i[j_2]$. If the following three conditions are satisfied, then the algorithm adds the minimum into the detected repetition set: 1) $r_i[j_1] - r_i[j_0] > h$ and $r_i[j_2] - r_i[j_0] > h$ (i.e., the local minimum is deep enough); 2) $(r_i[j_1] - r_i[j_0])/(j_1 - j_0) > p$ or $(r_i[j_2] - r_i[j_0])/(j_2 - j_0) > p$ (i.e., the local minimum is sharp enough); and 3) no two repetitions are closer than d.

Figure 2 shows the repetition detection result of a particular segment for the Beatles song "Yesterday." In Figure 2, the four detected local minima correspond to the four

		<i>V</i> ₁	<i>V</i> ₂	<i>V</i> ₃	 	 	 Vn
	0	0	0	0	 	 	 0
<i>u</i> ₁	е						
u ₂	2 _e						
u _m	me						

[FIG1] Dynamic time warping matrix WM with initial setting. *e* is a predefined parameter denoting the deletion cost.



[FIG2] One-segment repetition detection result of the Beatles song "Yesterday." The local minima indicated by circles correspond to detected repetitions of the segment.

restatements of the same melodic segment in the song ("Now it looks as though they are here to stay...," "There is a shadow hanging over me ...," "I need a place to hide away...," "I need a place to hide away"). However, the repetitions detected may have add or drop errors,

meaning a repetition is falsely detected or missed. The number of add and drop errors are balanced by the predefined parameter h; whenever the local minimum is deeper than

THE AMOUNT OF MUSIC AVAILABLE IN DIGITAL FORM HAS ALREADY BECOME UNFATHOMABLE.

height *h*, the algorithm reports a detection of repetition. Thus, when *h* increases, there are more drop errors but fewer add errors, and vise versa. For balancing between these two kinds of errors, the algorithm can search within a range for the best value of *h*, so that the number of detected repetitions of the whole song is reasonable (e.g., # total detected repetitions/ $n \approx 2$).

For each detected minimum $r_i[j^*]$ for S_i , let $k^* = t_i[j^*]$; thus, it is detected that segment $S_i = v_{k_i}v_{k_i+1}...v_{k_i+l-1}$ is repeated in V from $v_{k_i+k^*}$. Note that by the nature of dynamic programming, the matching part in V may not have length *l* due to the variations in the repetition.

SEGMENT MERGING

This step merges consecutive segments that have the same repetitive property into sections and generates pairs of similar sections. Figure 3 shows the repetition detection result of the Beatles song "Yesterday" after this step. In this figure, a circle or a square at (j, k) corresponds to a repetition detected in the last step (i.e., the segment starting from v_j is repeated from v_{j+k}). Since one musical phrase typically consists of multiple segments, based on the configurations in previous steps, if one segment in a phrase is repeated by a shift of k, all the segments in this phrase are repeated by shifts roughly equal to k.



[FIG3] Whole-song repetition detection result of the Beatles song "Yesterday."

This phenomenon can be seen from Figure 3, where the squares form horizontal patterns indicating that consecutive segments have roughly the same shifts.

By detecting these horizontal patterns (denoted by squares

in Figure 3) and discarding

AVAILABLE LREADY ABLE. other detected repetitions (denoted by circles in Figure 3), add or drop errors in repetition detection are further reduced. The output of this step is a set of sections consisting of

merged segments and the repetitive relation among these sections in terms of section-repetition vectors $[j_1 \ j_2 \ \text{shift}_1 \ \text{shift}_2]$, indicating that the segment starting from v_{j_1} and ending at v_{j_2} repeats roughly from $v_{j_1+\text{shift}_1}$ to $v_{j_2+\text{shift}_2}$. Each vector corresponds to one horizontal pattern in the whole-song repetition detection result. For example, the vector corresponding to the left-bottom horizontal pattern in Figure 3 is [200 520 370 370].

STRUCTURE LABELING

Based on the vectors obtained from the third step, the last step of the algorithm segments the entire piece into sections and labels each section according to the repetitive relation (i.e., gives each section a symbol such as "A," "B," etc.). This step will output the structural information, including both the form (e.g., AABABA) and the boundaries indicating the beginning and end of each section. To solve conflicts that might occur, the rule is to always label the most frequently repeated section first. Specifically, the algorithm finds the most frequently repeated section based on the first two columns in the section-repetition vectors and labels it and its shifted versions as section A. Then, the algorithm deletes the vector already labeled, repeats the same procedure for the remaining section-repetition vectors, and labels the sections produced in each step as B, C, D, and so on. If conflicts occur (e.g., a later labeled section has overlap with the previous labeled sections), the previously labeled sections will always remain intact, and the currently labeled section and its repetition will be truncated so that only the nonoverlapped part will be labeled as new.

MUSIC SUMMARIZATION

Music summarization (or thumbnailing) aims to find the most representative part of a musical piece. For example, for pop/rock songs, there are often catchy and repetitious parts (called the "hooks"), which can be implanted in your mind after hearing the song just once. This section analyzes the correlation between the representativeness of a musical part and its location within the global structure and proposes a method to automate music summarization. Results will be evaluated both by objective criteria and human experiments.

In general, it would be helpful for locating structurally accented locations (e.g., the beginning or the ending of a section, especially a chorus section) if the song has been segmented into meaningful sections before summarization. Once we have the recurrent structure of a song, we can have different music summarization strategies for different applications or different types of users. In the following, the methods we present will find the most representative part of music (specifically, hooks of pop/rock music) based on the result of recurrent structural analysis.

SECTION-BEGINNING STRATEGY (SBS)

The first strategy assumes that the most repeated part of the music is also the most representative part and that the beginning of a

section is typically essential. Thus, this strategy, illustrated in Figure 4, chooses the beginning of the most repeated section as the thumbnail of the music. The algorithm first finds the most repeated sections

SPECIFICALLY, GIVEN A MUSICAL PIECE (OR A PART OF IT), THE SYSTEM WILL SEGMENT IT INTO SECTIONS BASED ON KEY CHANGE AND IDENTIFY THE KEY OF EACH SECTION.

based on the structural analysis result, takes the first section among these, and truncates its beginning (20 s in this experiment) as the thumbnail.

SECTION-TRANSITION STRATEGY

We also investigated the music thumbnails at some commercial music Web sites for music sales (e.g., Amazon.com, music.msn.com) and found that the thumbnails they use do not always start from the beginning of a section and often contain the transition part (end of section A and beginning of section B). This strategy assumes that the transition part can give a good overview of both sections and is more likely to capture the hook (or title) of the song, though it typically will not give a thumbnail right at the beginning of a phrase or section.

Based on the structural analysis result, the algorithm finds a transition from section A to section B and then it truncates the end of section A, the bridge, and the beginning of section B



[FIG4] Section-beginning strategy.



[FIG5] Section-transition strategy.



[FIG6] An example for measuring segmentation performance: (a) detected transitions and (b) relevant transitions.

(shown in Figure 5). The boundary accuracy is not very important for this strategy.

To choose the transition for summarization, three methods were investigated:

• STS-I: Choose the transition such that the sum of the repeated times of A and of B is maximized; if there is more than one such transition, the first one will be chosen. In the above example, since there are only two different sections,

either $A \rightarrow B$ or $B \rightarrow A$ satisfies the condition. Thus, the first transition from A to B will be chosen.

• STS-II: Choose the most repeated transitions between different sections; if there is more than one such transi-

tion, the first one will be chosen. In the above example, $A \rightarrow B$ occurs twice, $B \rightarrow A$ occurs once; thus, the first transition from A to B will be chosen.

• STS-III: Choose the first transition right before the most repeated section. In the above example, B is the most repeated section; thus, the first transition from A to B will be chosen.

Although in the above example, all three methods will choose the same transition for summarization, we can come out with various other forms where the three methods will choose different transitions.

EXPERIMENT AND EVALUATION

EVALUATION OF SEGMENTATION

To evaluate segmentation results, two aspects need to be considered: label accuracy (whether the computed label of each frame



[FIG7] Detection of key change in "Mozart: Sonata No. 11 In A 'Rondo All Turca, 3rd movement'" (solid line: computed key; dotted line: truth). is consistent with the actual label) and segmentation accuracy (whether the detected locations of transitions are consistent with the actual locations).



[FIG8] Performance of key detection: (a) varying *stayprob* (W = 10) and (b) varying W (*stayprob* = 0.996).



[FIG9] Comparison of (a) the computed structure using DM and (b) the true structure of "Yesterday."

Label accuracy is defined as the proportion of frames that are labeled correctly, i.e.,

Label accuracy =
$$\frac{\text{#frames labeled correctly}}{\text{#total frames}}$$
. (11)

Two metrics were proposed and used for evaluating segmentation accuracy. Precision is defined as the proportion of detected transitions that are relevant. Recall is defined as the proportion of relevant transitions detected. Thus, if $B = \{\text{relevant transi$ $tions}\}$, $C = \{\text{detected transitions}\}$, and $A = B \cap C$, from the above definition, precision = A/C and recall = A/B.

To compute precision and recall, we need a parameter w: whenever a detected transition t_1 is close enough to a relevant transition t_2 such that $|t_1 - t_2| < w$, the transitions are deemed identical (a hit). Obviously, greater w will result in higher precision and recall. In the example shown in Figure 6, the width of each shaded area corresponds to 2w - 1. If a detected transition falls in a shaded area, there is a hit. Thus, the precision in this example is 3/6 = 0.5 and the recall is 3/4 = 0.75. Given w, higher precision and recall indicates better segmentation performance. In our experiment (512 window step at 11-kHz sampling rate), w will vary within a range to see how precision and recall vary accordingly: 10 frames (~0.46 s) to 80 frames (~3.72 s).

It can be shown that, given n and l, precision increases by increasing w (i.e., increasing m); recall increases by increasing k or w.

For recurrent structural analysis, besides label accuracy, precision, and recall, one extra metric—formal distance will be used to evaluate the difference between the computed form and the true form. It is defined as the edit distance between the strings representing different forms. For example, the formal dissimilarity between structure AABABA and structure AABBABBA is two, indicating two insertions from the first structure to the second structure (or, two deletions from the second structure to the first structure; thus, this definition of distance is symmetric). Note that how the system labels each section is not important as long as the repetitive relation is the same; thus, structure BBABAB is deemed as equivalent (0-distance) to structure BBABAB or structure AACACA.

EVALUATION OF THUMBNAILING

Based on the previous human experiments, five criteria for pop/rock music are considered for evaluating the summarization result. These criteria include: 1) the percentage of generated thumbnails that contain a vocal portion, 2) the percentage of generated thumbnails that contain the song's title, 3) the percentage of generated thumbnails that start at the beginning of a section, 4) the percentage of generated thumbnails that start at the beginning of a phrase, and 5) the percentage of generated thumbnails that capture a transition between different sections.

EXPERIMENTAL RESULTS

PERFORMANCE OF KEY DETECTION

Ten classical piano pieces were used in the experiment of key detection, since the chromagram representation of

piano music has a good mapping between its structure and its musical interpretation. These pieces were chosen randomly as long as they have fairly clear tonal structure (relatively tonal rather

THE CHROMAGRAM, ALSO CALLED THE PITCH CLASS PROFILE FEATURE, IS A FRAME-BASED REPRESENTATION OF AUDIO.

to the parts with no repetition.

tom two curves correspond to upper bounds of recall and precision based on random segmentation. The bottom horizontal line shows the baseline label accuracy of labeling the whole piece as one section.

The experimental result shows that the performance of seven out of ten piano pieces and 17 out of 26 Beatles songs have formal

the repetitions. Sections in the same color indicate restatements of the section. Sections in the lightest gray correspond

Figure 10 shows the segmentation performances of the two

data corpora, respectively, with varying w. In each plot, the bot-





than atonal). The "truth" was manually labeled by the author based on the score notation for comparison with the computed results. The data were mixed into 8-bit mono and down-sampled to 11 kHz. Each piece was segmented into

Figure 7 shows key detection results of Mozart's piano sonata No. 11 with stayprob = 0.996 in step 1 and stayprob 2 $= 1 - .1^{-20}$ in step 2. Figure 7(a) presents the result of key detection without considering mode (step 1) and Figure 7(b) presents the result of mode detection (step 2).

frames of 1,024 samples with an overlap of 512 samples.

To show that label accuracy, recall, and precision of key detection averaged over all the pieces, we can either fix wand change stayprob [Figure 8(a)], or fix stayprob and change w [Figure 8(b)]. In Figure 8(a), two groups of results are shown: one corresponds to the performance of step 1 without considering modes and the other corresponds to the overall performance of key detection taking mode into consideration. It clearly shows that when stayprob increases, precision also increases while recall and label accuracy decrease. In Figure 8(b), three groups of results are shown: one corresponds to the performance of step 1 without considering modes, one corresponds to the overall performance of key detection with mode taken into consideration, and one corresponds to recall and precision based on random segmentation.

Additionally, random label accuracy should be around 8%, without considering modes. It clearly shows that when w is increasing, the segmentation performance (recall and precision) is also increasing. Note that label accuracy is irrelevant to *w*.

PERFORMANCE OF

RECURRENT STRUCTURAL ANALYSIS

Two experimental corpora were tested. One corpus is piano music, which is the same as the one used for key detection. The other consists of the 26 Beatles songs in the two CD collection titled The Beatles (1962-1966). All of these musical pieces have clear recurrent structures, so that the true recurrent structures were labeled easily for comparison. The data were mixed into 8bit mono and down-sampled to 11 kHz.

To qualitatively evaluate the results, figures as shown in Figure 9 are used to compare the structure obtained from the algorithm to the true structure obtained by manually labeling

[FIG10] Segmentation performance of recurrent structural analysis. (a) Classical piano music and (b) Beatles songs.

[TABLE1] 20-S MUSIC SUMMARIZATION RESULT.

SBS STS-I STS-II	VOCAL 100% 96% 96%	TITLE 65% 73% 62%	BEGINNING OF A SECTION 62% 42% 31%	BEGINNING OF A PHRASE 54% 46% 46%	TRANSITION 23% 82% 91%
STS-III	96%	58%	31%	50%	82%

distances less than or equal to two. The label accuracy is significantly better than the baseline, and the segmentation performance is significantly better than random segmentation. This demonstrates the promise of the method.

We also found that the computed boundaries of each section were often slightly shifted from the true boundaries. This was mainly caused by the inaccuracy of the approximate pattern matching. To tackle this problem, other musical features (e.g., chord progressions, change in dynamics) should be used to detect local events so as to locate the boundaries accurately.

PERFORMANCE OF THUMBNAILING

Human experiments (not covered in this article) have shown that using the beginning of a piece is a fairly good summarization strategy for classical music. Here, we will only consider pop/rock music for evaluating summarization results. Table 1 shows the performance of all the strategies (SBS, STS-I, STS-II, and STS-III) presented in the "Music Summarization" section using the 26 Beatles songs. For evaluating transition criterion (5th column), only the 22 songs in our corpus that have different sections were counted.

The comparison of the thumbnailing strategies clearly shows that the section-transition strategies (STSs) generate a lower percentage of thumbnails starting at the beginning of a section or a phrase, while these thumbnails are more likely to contain transitions. SBS has the highest chance to capture the vocal, and STS-I has the highest chance of capturing the title. It is possible, though, to achieve better performance using this strategy if we can improve the structural analysis accuracy in the future.

CONCLUSIONS AND FUTURE WORK

This article presents our research into segmenting music based on its semantic structure (such as key change) and recurrent structure and summarizing music based on its structure. Experimental results were evaluated quantitatively, demonstrating the promise of the proposed methods. Future directions include inferring the hierarchical structures of music and incorporating more musical knowledge to achieve better accuracy. Furthermore, a successful solution to any of these problems depends on the study of human perception of music, for example, what makes part of music sound like a complete phrase and what makes it memorable or distinguishable. Human experiments are always necessary for exploring such questions.

AUTHOR

Wei Chai received the B.S. and M.S. degrees in computer science from Beijing University in 1996 and 1999, respectively. She received the M.S. and Ph.D. degrees from the MIT Media Laboratory in 2001 and 2005, respectively. Her dissertation research dealt with automatic analysis of musical structure for information retrieval. She has a wide range of interests in the application of machine learning, signal processing, and music cognition to audio and multimedia systems. She has been a research scientist at GE Global Research Center since 2005.

REFERENCES

[1] M.A. Bartsch and G.H. Wakefield, "To catch a chorus: Using chroma-based representations for audio thumbnailing," in *Proc. Workshop Applications of Signal Processing to Audio and Acoustics*, 2001.

[2] A.L. Berenzweig and D. Ellis, "Locating singing voice segments within music signals," in *Proc. Workshop Applications of Signal Processing to Audio and Acoustics*, NY, 2001.

[3] G. Burns, "A typology of 'hooks' in popular records," *Pop. Music*, vol. 6, pp. 1–20, Jan. 1987.

[4] W. Chai and B. Vercoe, "Music thumbnailing via structural analysis," in *Proc. ACM Multimedia Conf.*, 2003.

[5] W. Chai, "Structural analysis of musical signals via pattern matching," in *Proc. Int. Conf. Acoustics, Speech and Signal Processing*, 2003.

[6] W. Chai and B.L. Vercoe, "Structural analysis of musical signals for indexing and thumbnailing," in *Proc. Joint Conf. Digital Libraries*, 2003.

[7] C. Chuan and E. Chew, "Polyphonic audio key-finding using the spiral array CEG algorithm," in *Proc. Int. Conf. Multimedia and Expo*, Amsterdam, The Netherlands, July 6–8, 2005.

[8] R.B. Dannenberg and N. Hu, "Pattern discovery techniques for music audio," in Proc. Int. Conf. Music Information Retrieval, Oct. 2002.

[9] A. Sheh and D. Ellis, "Chord segmentation and recognition using EM-trained hidden Markov models," in *Proc. 4th Int. Symp. Music Information Retrieval ISMIR-03*, Baltimore, Oct. 2003.

[10] J. Foote, "Visualizing music and audio using self-similarity," in Proc. ACM Multimedia Conf., 1999.

[11] J. Foote, "Automatic audio segmentation using a measure of audio novelty," In *Proc. IEEE Int. Conf. Multimedia and Expo*, 2000, vol. WE, pp. 452–455.

[12] J.L. Hsu, C.C. Liu, and L.P. Chen, "Discovering nontrivial repeating patterns in music data," *IEEE Trans. Multimedia*, vol. 3, no. 3, pp. 311–325, Sept. 2001.

[13] B. Logan and S. Chu, "Music summarization using key phrases," in *Proc. Int. Conf. Acoustics, Speech and Signal Processing*, 2000.

[14] T. Kemp, M. Schmidt, M. Westphal, and A. Waibel, "Strategies for automatic segmentation audio data," in *Proc. Int. Conf. Acoustics, Speech and Signal Processing*, 2000.

[15] G. Peeters, A.L. Burthe, and X. Rodet, "Toward automatic music audio summary generation from signal analysis," in *Proc. Int. Conf. Music Information Retrieval*, Oct. 2002.

[16] L.R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, 1989.

[17] N. Scaringella, G. Zoia, and D. Mlynek, "Automatic genre classification of music content: A survey," *IEEE Signal Processing Mag.*, vol. 23, no. 2, pp. 133–141, 2006.

[18] C. Yang, "MACS: Music audio characteristic sequence indexing for similarity retrieval," in *Proc. Workshop Applications of Signal Processing to Audio and Acoustics*, 2001.