

# 15. Audio Source Separation in a Musical Context

Bryan Pardo, Zafar Rafii, Zhiyao Duan

When musical instruments are recorded in isolation, modern editing and mixing tools allow correction of small errors without requiring a group to re-record an entire passage. Isolated recording also allows rebalancing of levels between musicians without re-recording and application of audio effects to individual instruments. Many of these techniques require (nearly) isolated instrumental recordings to work. Unfortunately, there are many recording situations (e.g., a stereo recording of a 10-piece ensemble) where there are many more instruments than there are microphones, making many editing or remixing tasks difficult or impossible.

Audio source separation is the process of extracting individual sound sources (e.g., a single flute) from a mixture of sounds (e.g., a recording of a concert band using a single microphone). Effective source separation would allow application of editing and remixing techniques to existing recordings with multiple instruments on a single track.

In this chapter we will focus on a pair of source separation approaches designed to work with mu-

One of the great advances of the 20th and 21st centuries has been the introduction and refinement of audio recording and the audio editing techniques recording allows. Remixing, editing, and remastering in the studio has enabled the creation of new genres of music (e.g., *musique concrete*, modern hip hop). When instruments are recorded in isolation, modern editing and mixing tools allow correction of small errors in music recordings without requiring a group to re-record an entire passage (e.g., a single missed note in the flute part). This also allows rebalancing of levels between musicians without re-recording (e.g., increasing the volume of the flute compared to the trumpet) and application of audio effects to individual instruments (e.g., adding reverberation to the vocals, but not the bass).

15.1	<b>REPET</b> .....	286
15.1.1	Original REPET .....	286
15.1.2	Adaptive REPET .....	288
15.1.3	REPET-SIM .....	289
15.2	<b>Pitch-Based Source Separation</b> .....	291
15.2.1	Multipitch Estimation .....	291
15.2.2	Multipitch Streaming .....	292
15.2.3	Constructing Harmonic Masks .....	294
15.3	<b>Leveraging the Musical Score</b> .....	294
15.3.1	Audio-Score Alignment .....	294
15.3.2	Pitch Refinement and Source Separation .....	296
15.4	<b>Conclusions</b> .....	296
	<b>References</b> .....	297

sic audio. The first seeks the repeated elements in the musical scene and separates the repeating from the nonrepeating. The second looks for melodic elements, pitch tracking and streaming the audio into separate elements. Finally, we consider informing source separation with information from the musical score.

Many of these editing techniques require (nearly) isolated instrumental recordings. One cannot edit out a missed note in the flute, if that note is also recorded on the adjacent microphone for the violin. Unfortunately, there are many recording situations (e.g., a stereo recording of a 10-piece ensemble) where there are many more instruments than there are microphones. Thus, instruments are not recorded in isolation. Also, many legacy recordings are only available in the final stereo (two-channel) or mono (one-channel) mixture and the original premixed tracks are not available. All of these situations make many editing or remixing tasks difficult or impossible.

Audio source separation is the process of extracting individual sound sources (e.g., a single flute) from

a mixture of sounds (e.g., a recording of a concert band using a single microphone). Effective source separation would allow application of editing and remixing techniques to existing recordings with multiple instruments on a single track.

There have been many source separation approaches developed for general audio signals. Examples of general source separation algorithms include independent component analysis (ICA) [15.1], nonnegative matrix factorization (NMF) [15.2], nonnegative tensor factorization (NTF) [15.3], probabilistic latent component analysis (PLCA) [15.4], and robust principal component analysis (RPCA) [15.5]. While all of these have been applied to music, most do not leverage any particular features of music to aid in the separation process.

One element of a music scene which is highly salient is the repeating structure found in the music. *Schenker* asserted that repetition is what gives rise to the concept of the motive [15.6]. *Ruwet* used repetition as a criterion for dividing music into small parts, revealing the syntax of the musical piece [15.7]. *Ockelford* argued that repetition and imitation is what brings order to music [15.8].

Computational audio researchers have found repetitive elements in music audio useful for many purposes. Common applications include music summarization, e.g., see *Bartsch* [15.9], *Cooper* and *Foote* [15.10] and *Peeters* [15.11], audio segmentation [15.12], beat estimation [15.13], finding drum patterns in the audio [15.14], and structural analysis [15.15]. For a thorough review on music struc-

ture analysis, the reader is referred to [15.11, 16, 17].

The idea that repetition can be used for source separation is supported by recent findings in psychoacoustics. *McDermott* et al. established that the human auditory system is able to segregate individual sources by identifying them as repeating patterns embedded in the acoustic input, without requiring prior knowledge of the source properties [15.18]. Given this, could repetition in music be used as a basis for automatically separating the audio into a repeating background (e.g., a salsa montuno) and a varied foreground (the lead *salsero* sala singer)?

Another salient feature of a musical scene is melody. *Bregman* [15.19] shows that humans often link together distinct sound elements (e.g., notes) in time to produce perceptually salient entities called auditory streams. These streams often correlate strongly with melodies. Many trained musicians are able to segregate several melodies into independent elements they can attend to. Can a machine do the same? Often, a musical score can aid the trained musician to perform this task better. Can a score provide similar aid to an auditory streaming algorithm?

In this chapter we will focus on a pair of source separation approaches designed to work with music audio. The first seeks the repeated elements in the musical scene and separates the repeating from the non-repeating. The second looks for melodic elements, pitch tracking and streaming the audio into separate elements. This second approach is then informed by a musical score to improve performance.

## 15.1 REPET

In this work, we begin with the observation that passages in many kinds of folk and pop music can be understood as a background component that is generally repeating in time, with a superimposed foreground component that is generally variable in time (e.g., a repeating accompaniment superimposed with varying vocals or a solo instrument). On this basis the *repeating pattern extraction technique* (REPET) was proposed. REPET is an intuitive approach for separating the repeating background from the nonrepeating foreground in an audio mixture. The basic idea is to identify repeating elements in the mixture by measuring self-similarity along time, derive repeating models by averaging the repeating elements over their repetition rates, and extract the repeating structure by comparing the repeating models to the mixture.

A number of experiments have shown that REPET can be effectively applied to separate pop songs into their music accompaniment and singing voice. Unlike other approaches for source separation, REPET does not depend on special parametrizations, does not rely on complex frameworks, and does not require external information. Because it is only based on repetition, it has the advantage of being simple, fast, blind, and therefore completely and easily automatable. More information about REPET, including source code, audio examples, and related publications can be found at [15.20].

### 15.1.1 Original REPET

The original REPET algorithm was designed to separate the repeating background from the nonrepeating

foreground in an audio mixture (e.g., the music accompaniment from the singing voice in a pop song) by identifying a period and modeling a segment for the periodically repeating patterns [15.21, 22].

The method can be summarized in three stages (Fig. 15.1):

- Identification of a repeating period
- Modeling of a repeating segment
- Extraction of the repeating structure.

### Repeating Period Identification

In the first stage, the time-domain signal (top left of Fig. 15.1) is transformed into a time–frequency representation known as the spectrogram by using the short-time Fourier transform (STFT). A spectrogram represents sound as a two-dimensional structure, where the vertical axis shows frequency from low (at the bottom) to high. The horizontal axis shows time, from left to right. The spectrograms in Fig. 15.1 use red to indicate high energy and blue to show low energy.

The autocorrelation of each frequency channel is then computed. An autocorrelation measures the similarity of a signal with a delayed version of itself given different delays. The peak of the autocorrelation function indicates the period at which the sound repeats. To identify periodicity in the mixture, the beat spectrum [15.13] is derived from the spectrogram by

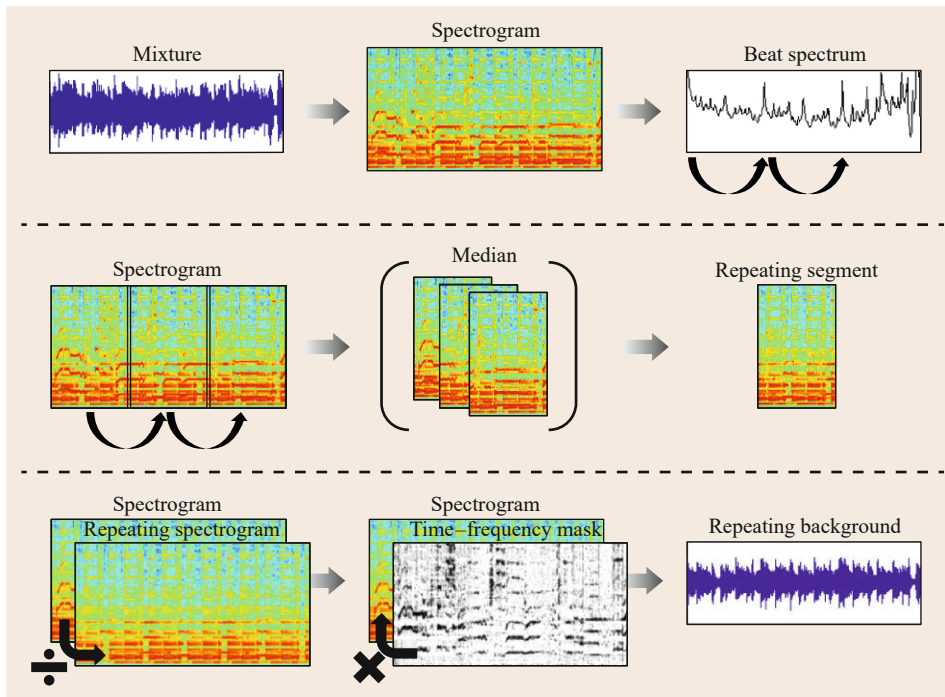
computing the autocorrelation over time for every frequency channel and averaging the autocorrelations over the frequency channels. This is shown in the upper right of Fig. 15.1. Here, a peak is a candidate period at which the audio may repeat.

If periodically repeating patterns are present in the mixture, the beat spectrum forms peaks that are periodically repeating at different period rates, unveiling the underlying periodically repeating structure of the mixture, as shown in Fig. 15.1. A best-fit repeating period can then be identified from the beat spectrum, manually or by using an automatic period finder [15.21, 22].

### Repeating Segment Modeling

In the second stage, the repeating period is used to find the period of the underlying repeating structure in the recording. This typically correlates to a pattern a few seconds long, such as a 4-chord repeating riff in a folk song. A model of what is repeating in the music is built by segmenting the audio at the points where the pattern repeats (middle panel of Fig. 15.1). The median value of the sound across all repetitions is then used to model the canonical repeating sound, as shown in Fig. 15.1.

This approach assumes the nonrepeating foreground (e.g., vocals or an instrumental soloist) has a sparse and varied time–frequency representation compared with the time–frequency representation of the



**Fig. 15.1** Overview of the original REPET: (1) computation of the beat spectrum and identification of a repeating period (*top row*); (2) filtering of the spectrogram and modeling of a repeating segment (*middle row*); (3) derivation of the time–frequency mask and extraction of the repeating background (*bottom row*)

repeating background. Therefore, time–frequency bins with small deviations at their repetition rate would most likely represent repeating elements and would be captured by the median model. Conversely, time–frequency bins with large deviations at their repetition rate would most likely be nonrepeating elements (i. e., the nonrepeating musical foreground) and would be removed by the median model.

### Repeating Structure Extraction

In the third stage, the repeating segment is used to derive a repeating spectrogram by taking, for every time–frequency bin, the minimum between the repeating model and the mixture spectrogram at period rate. This assumes the mixture spectrogram is the sum of a nonnegative repeating spectrogram and a nonnegative nonrepeating spectrogram. Thus, the mixture at any given time and frequency is assumed to always be as loud or louder than the individual mixture components (i. e., the repeating components are assumed not to cancel out the nonrepeating components).

The repeating spectrogram is then used to derive a time–frequency mask by dividing, for every time–frequency bin, the repeating spectrogram by the mixture spectrogram, as shown in Fig. 15.1. The rationale is that time–frequency bins that are likely to repeat at their repetition rate in the mixture spectrogram would have values near one in the time–frequency mask and would be weighted toward the repeating background. Conversely, time–frequency bins that are not likely to repeat at their repetition rate in the mixture spectrogram would have values near zero in the time–frequency mask and would be weighted toward the nonrepeating foreground.

The repeating background can then be obtained by multiplying, for every time–frequency bin, the time–frequency mask with the mixture. The nonrepeating foreground can be obtained by simply subtracting the repeating background from the mixture.

Experiments showed that REPET can be effectively applied to separate pop/rock song clips into their accompaniment music and singing voice [15.21, 22]. They also showed that REPET can be combined with other methods to improve background–foreground separation; for example, it can be used as a preprocessor to pitch detection algorithms to improve melody extraction [15.22], or as a postprocessor to a singing voice separation algorithm to improve music–voice separation [15.23]. Experiments further showed that REPET can be effectively applied to [separate the] accompaniment music and singing voice of separate full-track real-world songs, by simply applying the method along time via a sliding window [15.22]. There is, however, a trade-off for the window size in REPET: if the win-

dow is too long, the repetitions will not be sufficiently stable; if the window is too short, there will not be sufficient repetitions [15.22].

### 15.1.2 Adaptive REPET

Adaptive REPET is an extension of the original REPET and is designed to handle a varying periodic background, i. e., when the repeating period and/or the repeating patterns change over time (e.g., the succession of two homogenous sections in a song, such as a verse followed by the chorus), without the need for segmenting or windowing the audio [15.24].

As with the original REPET, the method can be summarized in three stages (Fig. 15.2):

- Identification of the repeating periods
- Modeling of a repeating spectrogram
- Extraction of the repeating structure.

#### Identification of Repeating Periods

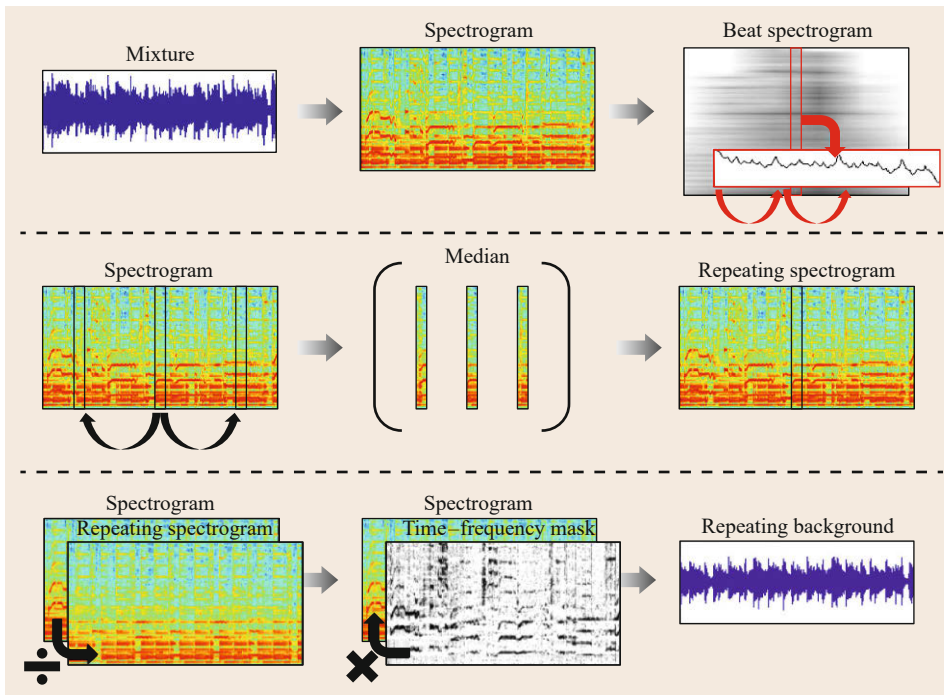
In the first stage, the signal is transformed into a spectrogram. To identify local periodicities in the mixture, the beat spectrogram [15.13] is derived from the spectrogram by computing a beat spectrum for every time frame by sliding a window along time. In other words, each column in the beat spectrogram represents a beat spectrum at a given time.

If periodically repeating patterns are present in the mixture, the beat spectrogram forms horizontal lines that are periodically repeating vertically, corresponding to the succession of peaks in the concatenated beat spectra, unveiling the underlying periodically repeating structure of the mixture, as shown in Fig. 15.2. If variations of periodicity happen over time in the mixture, the horizontal lines in the beat spectrogram will show variations in their vertical periodicity. The repeating periods can then be identified from the beat spectrogram, for all the time frames in the mixture spectrogram, manually or by using an automatic period finder [15.24].

#### Repeating Spectrogram Modeling

In the second stage, the repeating periods are used to model an initial repeating spectrogram by taking, for every time frame in the mixture spectrogram, the median of the time–frequency bins at their period rate, as shown in Fig. 15.2.

The original REPET assumes that there is a single large period (e.g., the length of a measure) for all the repeating elements in the audio. The adaptive REPET considers each time frame individually (a single frame lasts roughly 20–40 ms) and tries to find similar frames separated from the current frame by some period (e.g., 2 s). If similar frames are found at this period (e.g.,



**Fig. 15.2** Overview of adaptive REPET: (1) computation of the beat spectrogram and identification of the repeating periods (*top row*); (2) filtering of the spectrogram and modeling of a repeating spectrogram (*middle row*); (3) derivation of the time–frequency mask and extraction of the repeating background (*bottom row*)

similar frames at  $-4$  s,  $-2$  s,  $+2$  s), the repeating spectrogram model for the current frame is built from those frames. Once this is done, it moves forward to the next time frame and tries to find a period at which the content of the new frame repeats. This period may or may not be the same as the previous time frame (e.g., a period of 1.9 s instead of 2 s). This lets it handle periodically repeating structures that vary slowly over time and also structures that are composed of interleaved repeating patterns of different periods (e.g., minimalist music).

### Repeating Structure Extraction

In the third stage, the initial repeating spectrogram is used to derive a refined repeating spectrogram by taking, for every time–frequency bin, the minimum between the initial repeating spectrogram and the mixture spectrogram. The repeating spectrogram is then used to derive a time–frequency mask by dividing, for every time–frequency bin, the repeating spectrogram by the mixture spectrogram, as shown in Fig. 15.2.

The repeating background can then be obtained by multiplying, for every time–frequency bin, the time–frequency mask with the STFT of the mixture and transforming the result back to the time-domain. The nonrepeating foreground can be obtained by simply subtracting the repeating background from the mixture.

Experiments showed that adaptive REPET can be effectively applied to separate full-track real-world songs (e.g., a whole studio recording) into their accom-

paniment music and singing voice, unlike the original REPET which would only be meaningful for short excerpts [15.24].

### 15.1.3 REPET-SIM

REPET-SIM is a generalization of the REPET approach that was designed to handle nonperiodically repeating structures, i.e., when the repeating patterns happen intermittently or without a clear periodicity (e.g., repeated piano stabs in a jazz combo that use the same chord voicing, but whose rhythm varies). This is done by using a similarity matrix to identify the repeating elements [15.25].

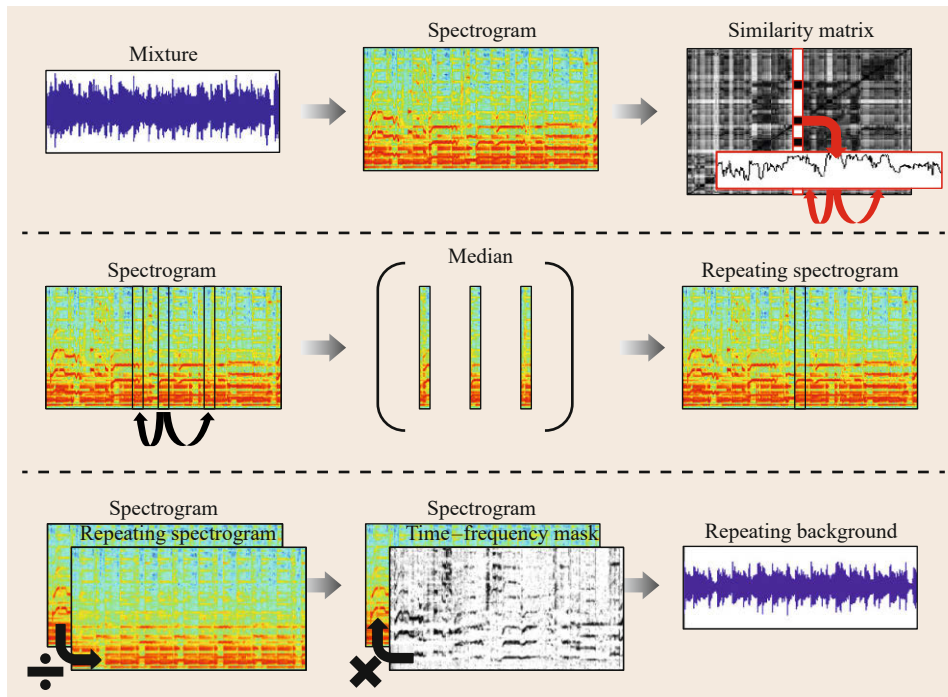
As with the previous two methods, the method can be summarized in three stages (Fig. 15.3):

- Identification of the repeating elements
- Modeling of a repeating spectrogram
- Extraction of the repeating structure.

#### Repeating Element Identification

In the first stage, the signal is transformed into a spectrogram. To identify similarities in the mixture, the similarity matrix [15.26] is derived from the spectrogram by computing the cosine similarity between any two pairs of time frames.

If repeating elements are present in the mixture, the similarity matrix would form regions of high and low



**Fig. 15.3** Overview of REPET-SIM: (1) computation of the similarity matrix and identification of the repeating elements (*top row*); (2) filtering of the spectrogram and modeling of a repeating spectrogram (*middle row*); (3) derivation of the time–frequency mask and extraction of the repeating background (*bottom row*)

similarity at different time indices, unveiling the underlying repeating structure of the mixture, as shown in Fig. 15.3. The repeating elements can then be identified from the similarity matrix, for all the time frames in the mixture spectrogram, manually or by using an automatic peak finder [15.25].

### Repeating Spectrogram Modeling

In the second stage, the repeating elements are used to model an initial repeating spectrogram by taking, for every time frame in the mixture spectrogram, the median of the time–frequency elements at their repetition rate, as shown in Fig. 15.3.

Compared with the other REPET methods (original and adaptive) that look for periodic similarity between events in the audio scene, REPET-SIM looks only for similarity, so that it can handle nonperiodically repeating structures, i. e., when the repeating patterns happen intermittently or without a clear periodicity.

### Repeating Structure Extraction

In the third stage, the initial repeating spectrogram is used to derive a refined repeating spectrogram by taking, for every time–frequency bin, the minimum be-

tween the initial repeating spectrogram and the mixture spectrogram.

The repeating spectrogram is then used to derive a time–frequency mask by dividing, for every time–frequency bin, the repeating spectrogram by the mixture spectrogram, as shown in Fig. 15.3.

The repeating background can then be obtained by multiplying, for every time–frequency bin, the time–frequency mask with the STFT of the mixture and transforming the result back to the time-domain. The nonrepeating foreground can be obtained by simply subtracting the repeating background from the mixture.

Experiments showed that REPET-SIM can be effectively applied to separate full-track real-world songs into their accompaniment music and singing voice, also compared with the adaptive REPET [15.25]. Experiments also showed that REPET-SIM can be effectively applied to separate two-channel mixtures of one speech source and real-world background noise into their background noise and clean speech, by applying the method online for real-time computing [15.27].

Note that *FitzGerald* proposed a method very similar to REPET-SIM for music–voice separation [15.28].

## 15.2 Pitch-Based Source Separation

The previous sections focused on performing source separation using cues related to the repetitive, rhythmic structure of the music audio. Another fruitful approach is to use the melodic content as embodied by the pitches present in the audio. We now turn to this approach.

Pitched musical instruments (e.g., brass, woodwinds, strings, and keyboard instruments) are harmonic sound sources. Most of the energy in a harmonic sound is located at frequencies that are integer multiples of the fundamental frequency (F0). For example, if a piano plays a single note at  $A = 440$  Hz, most of the energy in the sound will be concentrated at 440 Hz, 880 Hz, 1320 Hz, and so on. While F0 is a physical attribute and pitch is a perceptual attribute of a sound, for a harmonic sound, its pitch can be reliably matched to the F0. Therefore, we do not differentiate F0 from pitch in our discussions here.

Pitch information helps to separate harmonic sources from a mixture of sounds. In this section, we present our work on pitch-based source separation of audio mixtures composed of harmonic sound sources. Typically, this means separation of instruments from a music recording containing multiple concurrent instruments.

The first step is to estimate the pitches of these harmonic sources from the mixture [15.29]. This is called *multipitch estimation* (MPE). A frame is typically a 20–40 ms time window. MPE is typically performed in each individual time frame of the audio mixture. The second step is to connect the pitch estimates in different frames to form pitch trajectories, each of which corresponds to a source [15.30]. This is called *multipitch streaming*. The last step is to extract the harmonics from the pitch estimates of each source and reconstruct the source signal. In the following, we describe the three steps separately.

### 15.2.1 Multipitch Estimation

Multipitch estimation is the task of estimating pitches and the number of pitches in each frame of a harmonic sound mixture. In music information retrieval, the number of pitches is also called *polyphony*. Many methods have been proposed in the literature. Some do not employ any preprocessing of the signal and work with the full time domain or frequency domain signal [15.31–40]; others reduce the signal to a more compact representation such as employing an auditory filterbank as the front end [15.41–44] or representing the spectrum only with significant peaks [15.45–47]. Our method falls within the second category. More

specifically, we represent the power spectrum of the audio mixture with both significant peaks and nonpeak regions and propose a maximum likelihood method to estimate the pitches and the polyphony from the peak–nonpeak-region representation [15.29].

#### Peak Detection

For harmonic sounds, significant spectral peaks ideally correspond to harmonics of the pitches. Therefore, detection of these peaks would help us infer the pitches. Taking one frame of the audio signal, we first perform a Fourier transform [15.48] to turn the audio into a spectral representation. The middle top panel of Fig. 15.1 shows a spectrogram. Here, each column of the image is the spectral representation of a single time step (typically a 20–40 ms window).

Spectral peaks at each time slice are detected by the peak detector described in [15.49]. Basically, there are two criteria that determine whether a local maximum in the spectrum should be labeled as a peak. The first criterion is global: the local maximum should not be less than some threshold (e.g., 50 dB) lower than the maximum of the spectrum across all frequencies in that time step. The second criterion is local: the local maximum should be locally higher than a smoothed version of the spectrum by at least some threshold (e.g., 4 dB). Finally, the peak amplitudes and frequencies are refined by quadratic interpolation [15.50].

We define the *nonpeak region* as those frequencies that are further than a quarter tone from any of the detected spectral peaks. Although the nonpeak region cannot tell us where the pitches should be, it can tell us where the pitches should not be. Pitches are not likely to be located at frequencies whose harmonics are in the nonpeak region.

#### Likelihood Function

We propose a maximum likelihood method to estimate the set of pitches from the peak–nonpeak representation of the mixture spectrum. The likelihood function is defined as the multiplication of the *peak-region likelihood* and the *nonpeak-region likelihood*, assuming conditional independence of peaks and the nonpeak region given the pitches. The peak region likelihood is defined as the probability of occurrence of the peaks, given an assumed set of pitches. The nonpeak region likelihood is defined as the probability of *not* observing peaks in the nonpeak region, given an assumed set of F0s. The peak region likelihood and the nonpeak region likelihood act as a complementary pair. The former helps find F0s that have harmonics that explain peaks, while the latter helps avoid F0s that would predict har-

monics where none are observed (i. e., in the nonpeak region).

To define the peak-region likelihood, we characterize each peak by its frequency and amplitude. We further consider the probability that each detected peak is a *normal* peak or a *spurious* peak. By *normal* we mean the peak is generated by some harmonic of the underlying pitches; and by *spurious* we mean the peak is due to other factors such as side-lobes, peak detection errors, etc.

We train the parameters of the likelihood functions using training data with ground-truth pitches and detected peaks and nonpeak regions. Two kinds of training data are used. The first kind is a set of isolated notes from the University of Iowa Musical Instrument Samples (MIS) dataset [15.51]. They are used to train the parameters of the nonpeak region likelihood model. The second kind is a set of randomly mixed chords using these isolated notes. They are used to train the parameters of the peak-region likelihood model. Ground-truth pitches are detected using the YIN [15.52] pitch tracking algorithm on isolated notes before mixing the isolated notes together. Peaks and the nonpeak region are detected using the proposed method.

#### Pitch and Polyphony Estimation

Given the set of detected peaks, frequencies within one semitone around each peak are considered as possible pitches (i. e., possible F0s). The underlying pitches in this frame are thus assumed to compose a subset of these frequencies. This helps to constrain the set of possible hypotheses to reasonable ones, given the data. The task of the maximum likelihood estimation is thus to find the subset of potential F0s that gives the highest likelihood to having generated the observed peaks in the spectrum.

It is, however, intractable to enumerate all these subsets. A typical scene may have 50–100 peaks. The set of all subsets of 100 peaks has  $2^{100}$  elements and is not tractable to fully search. We therefore propose an iterative greedy search strategy [15.29]. We start from an empty subset to represent the estimated pitches. At each iteration, we add the one pitch estimate that most increases the likelihood of the subset. This strategy only enumerates a very small amount of subsets, as the choice of latter pitches depends on the choice of earlier pitches. However, the computational complexity is significantly reduced from exponential to linear with respect to the number of peaks.

An important question for this iterative greedy search process is *When should we stop?* In other words, how many pitches should we estimate in each time frame. Ideally, we hope that the likelihood function can take care of this problem and stop the process when the

likelihood no longer increases on adding a new pitch. However, similar to many other maximum likelihood estimation problems, there is an overfitting problem. The likelihood typically increases with the number of pitches, although the increase becomes slower. We use a simple thresholding method to address this problem. We do not stop the iterative process until the number of pitches in the subset reaches a predefined maximally possible polyphony (e.g., 9). We then calculate the likelihood increase from 1 pitch to the maximally possible polyphony as the maximally possible likelihood increase. We then estimate the polyphony as the number of pitches that first surpasses 88% of the maximally possible increase. The threshold was tuned on a training set of musical chords with polyphony from 1–6. This simple polyphony estimation method is shown to work well on both musical chords and real music pieces.

#### Pitch Refinement

Pitches and polyphony estimated in individual frames may contain errors that make the pitch contours non-smooth over time. By utilizing contextual information, it is possible to fix these errors and improve the pitch estimation accuracy.

We use a moving average to calculate the refined polyphony estimate with a triangular moving window with size of 19 frames. We also build a pitch histogram with a granularity of a semitone by counting the pitch estimates within the window and rank the pitches by their weighted counts according to the window. The top several pitches are returned as the refined pitches with equal weights, where the number is equal to the refined polyphony estimate. In experiments we show that this refinement step removes many insertion and deletion pitch estimation errors and significantly increases the estimation accuracy [15.29].

### 15.2.2 Multipitch Streaming

The pitches estimated in the previous section can help us separate harmonic sources in each individual frame. However, to separate the source signal over multiple frames, we need to connect these pitches into streams, each of which corresponds to a source. This is the multipitch streaming problem. Researchers have proposed using frequency–amplitude continuity to stream pitches [15.31, 36, 53, 54], but this approach only works within a note where the pitch does not change abruptly. For streaming pitches into noncontinuous pitch contours, we proposed the first method [15.30]. The basic idea is to use the timbre information. Notes performed by the same instrument have similar timbre compared to those performed by different instruments. Therefore,



we associate each pitch estimate with a timbre feature vector and perform clustering on the timbre feature vectors. Ideally, each cluster will correspond to one source and their pitches form the pitch stream of that source.

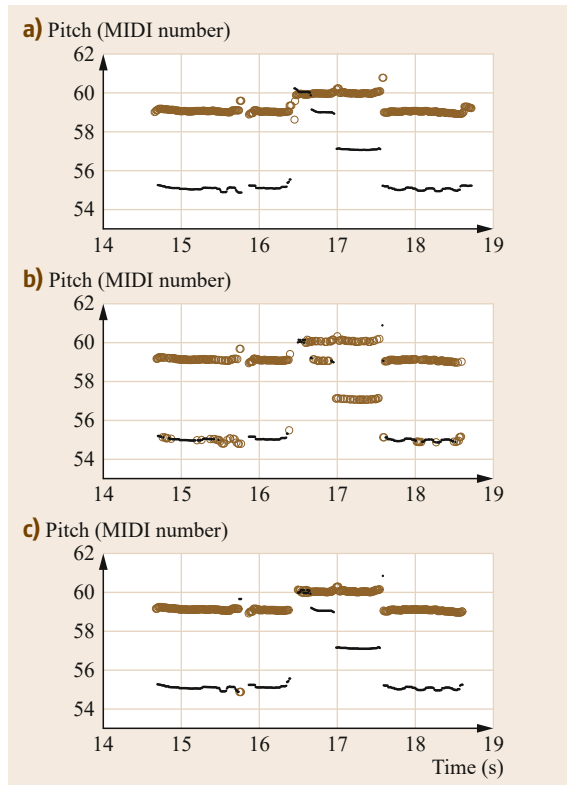
### Timbre Features

We need to calculate a timbre feature vector for each pitch estimate in each frame and this feature should be calculated from the mixture signal directly. In our work we use two kinds of features. The first is called *harmonic structure* and is described in [15.49]. It is defined as a vector of relative logarithmic amplitudes of the harmonics of a pitch estimate. The harmonics are at integer multiples of the pitch. We use the first 50 harmonics to create a 50-dimensional timbre vector. We choose this dimensionality because most instruments have less than 50 prominent harmonics. For each harmonic, we use the peak-finder from [15.49] to see if there is a significant peak within a musical quarter-tone. If no peak is associated, the magnitude of the harmonic is set to 0 dB, otherwise it is set to the value of the nearest peak. Then, the representation is normalized. This feature has been shown to be similar for notes played by the same instrument within a narrow pitch range, while it is different for different instruments.

Another feature is called the *uniform discrete cepstrum* (UDC) [15.30]. It is calculated by taking the discrete cosine transform of a sparse log-amplitude magnitude spectrum where the nonzero elements are the harmonics of the pitch. This feature lies in the cepstral feature category and represents the spectral envelope of the harmonics of the target pitch. However, unlike other cepstral features such as the ordinary cepstrum or mel-frequency cepstral coefficients (MFCC), UDC can be calculated from the mixture spectrum directly for the target source, without resorting to source separation. UDC has been shown to outperform other cepstral representations and the harmonic structure feature in an instrument recognition task of musical chords [15.55].

### Constrained Clustering

Given the feature vector of each pitch estimate, we perform  $K$ -means clustering on the pitches, to minimize the timbre inconsistency within each cluster. However, results show that there are two kinds of common errors. In Fig. 15.4b, a number of pitches are clustered into the wrong trajectory. For example, the pitches around Musical Instrument Digital Interface (MIDI) number 55 from 14.8 to 15.8 s form a continuous contour and are all played by the bassoon. However, in the clustering, some of them are assigned to the saxophone. In another example, from 16.8 to 17.6 s, the  $K$ -means



**Fig. 15.4a–c** Comparison of the ground-truth pitch streams (a),  $K$ -means clustering ( $K = 2$ ) results (i. e., only minimizing the objective function) (b) and the proposed method’s results (i. e., considering both objective and constraints) (c). Both the  $K$ -means and the proposed method take the ground-truth pitches as inputs, use 50 d harmonic structure as the timbre feature and randomly initialize their clusterings. Each point in these figures is a pitch. Different instruments are marked with different markers (circles for saxophone and dots for bassoon)

clustering places two simultaneous pitches into the saxophone stream. This is not reasonable as the saxophone is a monophonic instrument.

If we know that different sources do not often perform the same pitch at the same time and all sources are monophonic, we can impose two kinds of constraints on some pairs of the pitches to improve clustering: A *must-link* constraint is imposed between two pitches that differ less than  $\Delta_t$  in time and  $\Delta_f$  in frequency. It specifies that two pitches close in both time and frequency should be assigned to the same cluster. A *cannot-link* constraint is imposed between two pitches in the same frame. It specifies that two simultaneous pitches should be assigned to different clusters. These must-links and cannot-links form the set of all constraints. Figure 15.4c shows the result obtained from our pro-

posed algorithm, considering both the objective and constraints.

### Iterative Algorithm

The formulated constrained clustering problem has two properties that mean existing constrained clustering algorithms [15.56–58] cannot be applied: 1) constraints are inconsistent with each other as they are imposed on pitch estimates that contain errors; 2) almost every pitch estimate is involved in some constraint. We therefore propose a new algorithm. It starts from an initial partition that satisfies a subset of all the constraints. Then it iteratively minimizes the objective function while incrementally satisfying more constraints. While the details of the algorithm can be found in [15.30], here we state its two important properties: 1) it always converges; 2) the timbre inconsistency objective function monotonically strictly decreases while the number of satisfied constraints monotonically increases.

### 15.2.3 Constructing Harmonic Masks

Given the estimated pitch stream for each harmonic source, we build a soft frequency mask around its harmonics to separate its magnitude spectrum from the mixture spectrum [15.59]. We then combine the separated magnitude spectrum with the phase spectrum of the mixture signal and perform an inverse Fourier transform to calculate the time-domain signal. Finally, the overlap-add technique [15.48] is applied to concatenate the current frame to previously separated frames.

## 15.3 Leveraging the Musical Score

Previous sections described source separation algorithms that depend on repeating (rhythmic) content and pitch-based (melodic) content to perform separation. We now consider the case where one can improve a source separation algorithm by using information in addition to the audio recording. For many music pieces, music scores are widely available. In this scenario, score information can be leveraged to help analyze and separate the audio signal. More specifically, if the audio performance is faithful to the score and the audio and the score are aligned (i. e., synchronized), the score can tell us what pitches are likely being played at a time of the audio. This can greatly improve the accuracy of melodic, pitch-based source separation. We can use the score-provided pitch information to separate the harmonic sources in the audio. Such a system is called a score-informed source separation system.

To calculate the frequency masks for sources, we first identify their harmonics and overlapping situations from the estimated pitches. Each frequency bin is then classified into three kinds according to the number of harmonics that involve this frequency bin:

- Nonharmonic bin
- Nonoverlapping harmonic bin
- Overlapping harmonic bin.

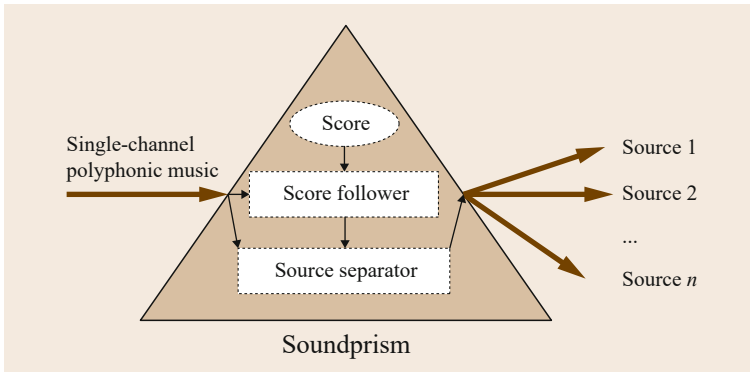
For a nonharmonic bin, the masks are designed to evenly distribute the mixture energy to all active sources. For a nonoverlapping harmonic bin, the mixture energy is solely distributed to the source whose harmonic involves the bin.

The mask design for overlapping harmonic bins is the most difficult. One can calculate an average harmonic structure template for each source from the harmonic structures of its estimated pitches and then use the template to design the mask value at different harmonics [15.49]. This method considers the timbre model of the sources. A simpler method is to distribute the mixture energy to overlapping harmonics, in the inverse proportion to the square of the harmonic indices. This method does not model the timbre of sources, instead, it makes a general assumption that the harmonic amplitude decays at the same rate with respect to the harmonic index, regardless of pitch and instrument that produced the note. This is a very coarse assumption but it provides a simple and relatively effective way to resolve overlapping harmonics.

In this chapter, we describe our work on an online score-informed source separation system called *Soundprism* [15.59]. Soundprism first aligns the audio with the score, then estimates the precise pitches based on the score-provided pitches in each frame and finally separates audio sources in the frame. All of these operations are performed in an online fashion, i. e., it processes the audio frame-by-frame in a serial fashion, without having the entire audio available. Figure 15.5 illustrates the system overview of Soundprism.

### 15.3.1 Audio-Score Alignment

The first stage of Soundprism is online audio-score alignment, also called *score following*. The score follower takes a piece of polyphonic music audio and its electronic score (e.g., MIDI) as input. It then outputs a score position for each time frame in a sequence. We




**Fig. 15.5** System overview of Soundprism

propose a hidden Markov process model to do so, as illustrated in Fig. 15.6. The  $n$ -th time frame of the audio is associated with a 2-dimensional hidden state vector  $\mathbf{s}_n = (x_n, v_n)^\top$ , where  $x_n$  is its score position (in beats) and  $v_n$  is its tempo (in beats per minute [BPM]). Each audio frame is also associated with an observation variable  $\mathbf{y}_n$ , which represents the magnitude spectrum of the time frame. Our aim is to infer the current score position  $x_n$  from current and previous observations  $\mathbf{y}_1, \dots, \mathbf{y}_n$ . To do so, we need to define a process model to describe how the states transition, an observation model to evaluate hypothesized score positions for the current audio frame, and to find a way to do the inference in an online fashion.

### Process Model

A process model defines the transition probability from the previous state to the current state, i. e.,  $p(\mathbf{s}_n | \mathbf{s}_{n-1})$ . This tells us how the score position and tempo changes



**Fig. 15.6** Illustration of the state space model for online audio-score alignment

from one frame to another and the probability of the change. We use two dynamic equations to define this transition. While the detailed equations are ignored here, the basic idea is that the change of score position is determined by the tempo and the time interval between two adjacent frames. Also the tempo changes randomly around the previous tempo assuming a Gaussian distribution if the score position just passed a note onset or offset and does not change otherwise.

Note that we do not introduce randomness directly in score position. This is to avoid disruptive changes of score position estimates. In addition, randomness is only introduced when the score position has just passed a note onset or offset. This is because it is rather rare that the performer changes tempo within a note. Second, on the listener's side, it is impossible to detect the tempo change before hearing an onset or offset, even if the performer does make a change within a note. Therefore, changing the tempo state in the middle of a note is not in accordance with music performance, nor does it have evidence to estimate this change.

### Observation Model

The observation model evaluates how well a hypothesized state (score position and tempo) can explain the observation, i. e.,  $p(\mathbf{y}_n | \mathbf{s}_n)$ . In this work, we use the multipitch likelihood model as described in Sect. 15.2.1. The basic idea is that if the score pitches at the hypothesized score position fit well to the magnitude spectrum of the current frame  $\mathbf{y}_n$ , then the hypothesis is good. To calculate the multipitch likelihood, we just plug the score pitches into the likelihood function described in Sect. 15.2.1.

Clearly the observation model itself is not enough to estimate the hidden state (e.g., the score position), as the score may show the same pitches at different positions. In addition, the tempo dimension of the state does not play in the observation model. These problems, however, are addressed when the observation model and process model work together. The process model, con-

sidering the tempo dimension and the continuity of state changes, would only favor hypothesized states whose score position is close to the previous score position. This is the key idea of the hidden Markov process model.

Our observation model only considers information from the current frame and could be improved if considering information from multiple frames. Ewert et al. [15.60] incorporate interframe features to utilize note onset information and improve the alignment accuracy. Joder et al. [15.61] propose an observation model which uses observations from multiple frames for their conditional random field-based method. In the future we want to explore these directions to improve our score follower.

### Inference

Given the process model and the observation model, we want to infer the state of the current frame from current and past observations. From a Bayesian point of view, this means we first estimate the posterior probability  $p(s_n | y_1, \dots, y_n)$ , then decide its value using some criterion like maximum a posteriori (MAP) or minimum mean square error (MMSE). For hidden Markov processes, this posterior probability at the current frame can be updated from the previous frame. Therefore, we can estimate the posterior probability and the hidden states in an online fashion.

Here we use a *bootstrap filter*, one variant of particle filters [15.62, 63] to do the online update of the posterior probability. The process starts from an initialization of  $M$  particles. Their score positions are all set to the beginning of the score and their tempi are uniformly distributed between half and twice of the score-notated tempo. These particles represent an initialization of the posterior probability. To update the posterior probability when a new audio frame comes in, the particles are first moved using the process model, i. e., the score positions and tempi of the particles are changed. Then the observation likelihood of each particle is calculated using the observation model, which indicates the fitness of the particle to the current audio frame. The likelihood is set as the weight of the particle. These particles are then resampled with replacement according to their weights to generate a new set of  $M$  particles. Particles that do not fit to the current frame are less likely to remain in

## 15.4 Conclusions

In this chapter we have outlined how to perform audio source separation on music using the cues of repeating musical structure and pitch content. We then

the new particle set, while particles that are a better fit to the current frame may retain multiple copies in the new particle set. A small random perturbation is imposed on these copies to prevent degeneracy of the particles. Now the new set of particles represent the new posterior probability. The average value of these particles is output as the estimate of the hidden state in the current frame.

The set of particles is not able to represent the distribution if there are too few and is time-consuming to update if there are too many. In our work we tried to use 100, 1000, and 10 000 particles. We find that with 100 particles, the score follower is often lost after a number of frames. But with 1000 particles, this rarely happens and the update is still fast enough. Therefore, 1000 particles are used in this chapter.

### 15.3.2 Pitch Refinement and Source Separation

Given the aligned score position for each audio frame, we know what instrument is playing what pitch in this frame from the score. This is important information for separating harmonic sources. However, the pitches provided by the score are integer MIDI pitch numbers. MIDI pitch numbers indicate keys on the piano keyboard. Typically, MIDI 69 indicates the A above Middle C. Assuming A440-based equal temperament allows translation from MIDI pitch to frequency in Hz. The resulting frequencies are rarely equal to the real pitches played in an audio performance. In order to extract the harmonics of each source in the audio mixture, we need to refine them to get accurate estimates of pitches played in the audio.

We refine the pitches using the multipitch estimation algorithm as described in Sect. 15.2.1, but restricting the search space within a semitone of the score-notated pitches. We also assume polyphony is given by the score.

Given the refined pitches in each audio frame, we use the same method as described in Sect. 15.2.3 to construct a harmonic mask for each pitch to separate the magnitude spectrum. Finally, we apply the inverse Fourier transform with the phase spectrum of the mixture signal and the overlap-add technique to reconstruct the separated source signals.

showed how one can augment the pitch-based separation algorithm by leveraging the information in a musical score, where available. Moving forward,

we envision a combination of the score-informed pitch-based separation with separation based on rhythmic structure. Combining these should allow separa-

tion of musical instruments from an audio mixture in a large variety of contexts that are currently intractable.

## References

- 15.1 P. Common, C. Jutten (Eds.): *Handbook of Blind Source Separation: Independent Component Analysis and Applications*, 1st edn. (Academic, Oxford 2010)
- 15.2 T. Virtanen: Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria, *IEEE Trans. Audio Speech Lang. Process.* **15**(3), 1066–1074 (2007)
- 15.3 D. FitzGerald, M. Cranitch, E. Coyle: Non-negative tensor factorisation for sound source separation. In: *Irish Signals and Syst. Conf., Dublin* (2005)
- 15.4 P. Smaragdis, B. Raj, M.V.S. Shashanka: A probabilistic latent variable model for acoustic modeling. In: *NIPS Workshop Adv. Modeling Acoust. Process., Whistler* (2006)
- 15.5 P.-S. Huang, S.D. Chen, P. Smaragdis: Singing-voice separation from monaural recordings using robust principal component analysis. In: *37th Int. Conf. Acoustics, Speech and Signal Process., Kyoto* (2012)
- 15.6 H. Schenker: *Harmony*, Vol. 1 (Univ. Chicago Press, Chicago 1980)
- 15.7 N. Ruwet, M. Everist: Methods of analysis in musicology, *Music Anal.* **6**(1/2), 3–9 (1987)
- 15.8 A. Ockelford: *Repetition in Music: Theoretical and Metatheoretical Perspectives*, Royal Musical Association Monographs, Vol. 13, 2005)
- 15.9 M.A. Bartsch: To catch a chorus using chroma-based representations for audio thumbnailing. In: *IEEE Workshop Appl. Signal Process. Audio Acoust., New Paltz* (2001)
- 15.10 M. Cooper, J. Foote: Automatic music summarization via similarity analysis. In: *3rd Int. Conf. Music Inf. Retr., Paris* (2002)
- 15.11 G. Peeters: Deriving musical structures from signal analysis for music audio summary generation: Sequence and state approach, *Comput. Music Modeling Retr.* **2771**, 143–166 (2004)
- 15.12 J. Foote: Automatic audio segmentation using a measure of audio novelty. In: *IEEE Int. Conf. Multimedia and Expo, New York* (2000)
- 15.13 J. Foote, S. Uchihashi: The beat spectrum: A new approach to rhythm analysis. In: *IEEE Int. Conf. Multimedia and Expo, Tokyo* (2001)
- 15.14 K. Yoshii, M. Goto, H.G. Okuno: Drum sound identification for polyphonic music using template adaptation and matching methods. In: *ISCA Tutor. Res. Workshop on Stat. Percept. Audio Process., Jeju* (2004)
- 15.15 R.B. Dannenberg: Listening to Naima: An automated structural analysis of music from recorded audio. In: *Int. Comput. Music Conf., Gothenburg* (2002)
- 15.16 R.B. Dannenberg, M. Goto: Music structure analysis from acoustic signals, *Handbook of Signal Process. Acoustics* **1**, 305–331 (2009)
- 15.17 J. Paulus, M. Müller, A. Klapuri: Audio-based music structure analysis. In: *11th Int. Soc. Music Inf. Retr., Utrecht* (2010)
- 15.18 J.H. McDermott, D. Wroblewski, A.J. Oxenham: Recovering sound sources from embedded repetition, *Proc. Nat. Acad. Sci. USA* **108**(3), 1188–1193 (2011)
- 15.19 A. Bregman, C. Jutten: *Auditory Scene Analysis: The Perceptual Organization of Sound* (MIT Press, Cambridge 1994)
- 15.20 Interactive Audio Lab of Northwestern University: <http://music.eecs.northwestern.edu/research.php?project=repet>
- 15.21 Z. Rafii, B. Pardo: A simple music-voice separation system based on the extraction of the repeating musical structure. In: *36th Int. Conf. Acoust. Speech Signal Process., Prague* (2011)
- 15.22 Z. Rafii, B. Pardo: REpeating pattern extraction technique (REPET): A simple method for music-voice separation, *IEEE Trans. Audio Speech Lang. Process.* **21**(1), 71–82 (2013)
- 15.23 Z. Rafii, D.L. Sun, F.G. Germain, G.J. Mysore: Combining modeling of singing voice and background music for automatic separation of musical mixtures. In: *14th Int. Soc. Music Inf. Retr., Curitiba* (2013)
- 15.24 A. Liutkus, Z. Rafii, R. Badeau, B. Pardo, G. Richard: Adaptive filtering for music-voice separation exploiting the repeating musical structure. In: *37th Int. Conf. Acoustics, Speech and Signal Process., Kyoto* (2012)
- 15.25 Z. Rafii, B. Pardo: Music-voice separation using the similarity matrix. In: *13th Int. Soc. Music Inf. Retr., Porto* (2012)
- 15.26 J. Foote: Visualizing music and audio using self-similarity. In: *7th ACM Int. Conf. Multimedia, Orlando* (1999)
- 15.27 Z. Rafii, B. Pardo: Online REPET-SIM for real-time speech enhancement. In: *38th Int. Conf. Acoust. Speech and Signal Process., Vancouver* (2013)
- 15.28 D. FitzGerald: Vocal separation using nearest neighbours and median filtering. In: *23rd IET Irish Signals and Syst. Conf., Maynooth* (2012)
- 15.29 Z. Duan, B. Pardo, C. Zhang: Multiple fundamental frequency estimation by modeling spectral peaks and non-peak regions, *IEEE Trans. Audio Speech Lang. Process.* **18**(8), 2121–2133 (2010)
- 15.30 Z. Duan, J. Han, B. Pardo: Multi-pitch streaming of harmonic sound mixtures, *IEEE Trans. Audio Speech Lang. Process.* **22**(1), 1–13 (2014)
- 15.31 G.E. Poliner, D.P.W. Ellis: A discriminative model for polyphonic piano transcription, *EURASIP J. Adv. Signal Process.* **2007**, 48317–1–48317–9 (2007), <https://doi.org/10.1155/2007/48317>

- 15.32 M. Davy, S.J. Godsill, J. Idier: Bayesian analysis of polyphonic western tonal music, *J. Acoustical Soc. Am.* **119**, 2498–2517 (2006)
- 15.33 E. Vincent, M.D. Plumbley: Efficient Bayesian inference for harmonic models via adaptive posterior factorization, *Neurocomputing* **72**, 79–87 (2008)
- 15.34 K. Kashino, H. Murase: A sound source identification system for ensemble music based on template adaptation and music stream extraction, *Speech Commun.* **27**(3–4), 337–349 (1999)
- 15.35 M. Goto: A real-time music-scene-description system: Predominant-F0 estimation for detecting melody and bass lines in real-world audio signals, *Speech Commun.* **43**(4), 311–329 (2004)
- 15.36 H. Kameoka, T. Nishimoto, S. Sagayama: A multipitch analyzer based on harmonic temporal structured clustering, *IEEE Trans. Audio Speech Lang. Process.* **15**(3), 982–994 (2007)
- 15.37 S. Saito, H. Kameoka, K. Takahashi, T. Nishimoto, S. Sagayama: Spectral analysis of polyphonic music signals, *IEEE Trans. Speech Audio Process.* **16**(3), 639–650 (2008)
- 15.38 J.-L. Durrieu, G. Richard, B. David: Singer melody extraction in polyphonic signals using source separation methods. In: *Proc. IEEE Int. Conf. Acoustics Speech Signal Process. (ICASSP)* (2008) pp. 169–172
- 15.39 V. Emiya, R. Badeau, B. David: Multipitch estimation of quasi-harmonic sounds in colored noise. In: *Proc. Int. Conf. Digital Audio Effects (DAFx)* (2007)
- 15.40 G. Reis, N. Fonseca, F. Ferndandez: Genetic algorithm approach to polyphonic music transcription. In: *Proc. IEEE Int. Symp. Intell. Signal Process* (2007)
- 15.41 T. Tolonen, M. Karjalainen: A computationally efficient multipitch analysis model, *IEEE Trans. Speech Audio Process.* **8**(6), 708–716 (2000)
- 15.42 A. de Cheveigné, H. Kawahara: Multiple period estimation and pitch perception model, *Speech Commun.* **27**, 175–185 (1999)
- 15.43 A. Klapuri: Multiple fundamental frequency estimation based on harmonicity and spectral smoothness, *IEEE Trans. Speech Audio Process.* **11**(6), 804–815 (2003)
- 15.44 A. Klapuri: Multiple fundamental frequency estimation by summing harmonic amplitudes. In: *Proc. ISMIR* (2006) pp. 216–221
- 15.45 R.J. Leistikow, H.D. Thornburg, J.S. Smith, J. Berger: Bayesian identification of closely-spaced chords from single-frame STFT peaks. In: *Proc. Int. Conf. Digital Audio Effects (DAFx'04), Naples* (2004) pp. 228–233
- 15.46 A. Pertusa, J.M. Inesta: Multiple fundamental frequency estimation using Gaussian smoothness. In: *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)* (2008) pp. 105–108
- 15.47 C. Yeh, A. Röbel, X. Rodet: Multiple fundamental frequency estimation of polyphonic music signals. In: *Proc. IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)* (2005) pp. 225–228
- 15.48 J.O. Smith: Spectral Audio Signal Processing, <http://ccrma.stanford.edu/~jos/saspl/> (2014)
- 15.49 Z. Duan, Y. Zhang, C. Zhang, Z. Shi: Unsupervised single-channel music source separation by average harmonic structure modeling, *IEEE Trans. Audio Speech Lang. Process.* **16**(4), 766–778 (2008)
- 15.50 J.O. Smith, X. Serra: Parshl: An analysis–synthesis program for non-harmonic sounds based on a sinusoidal representation. In: *Proc. Int. Comput. Music Conf. (ICMC)* (1987)
- 15.51 L. Fritts, University of Iowa: <http://theremin.music.uiowa.edu/MIS.html>
- 15.52 A. de Cheveigné, H. Kawahara: YIN, a fundamental frequency estimator for speech and music, *J. Acoustical Soc. Am.* **111**, 1917–1930 (2002)
- 15.53 M. Ryyanen, A. Klapuri: Polyphonic music transcription using note event modeling. In: *Proc. IEEE Workshop on Appl. Signal Process. Audio Acoustics (WASPAA)* (2005) pp. 319–322
- 15.54 W.-C. Chang, A.W.Y. Su, C. Yeh, A. Robel, X. Rodet: Multiple-F0 tracking based on a high-order HMM model. In: *Proc. Int. Conf. Digital Audio Effects (DAFx)* (2008)
- 15.55 Z. Duan, B. Pardo, L. Daudet: A novel cepstral representation for timbre modeling of sound sources in polyphonic mixtures. In: *Proc. IEEE Int. Conf. Acoustics Speech Signal Process. (ICASSP)* (2014)
- 15.56 K. Wagstaff, C. Cardie: Clustering with instance-level constraints. In: *Proc. Int. Conf. Machine Learning (ICML)* (2000) pp. 1103–1110
- 15.57 K. Wagstaff, C. Cardie, S. Rogers, S. Schroedl: Constrained k-means clustering with background knowledge. In: *Proc. Int. Conf. Machine Learning (ICML)* (2001) pp. 577–584
- 15.58 I. Davidson, S.S. Ravi, M. Ester: Efficient incremental constrained clustering. In: *Proc. ACM Conf. Knowl. Discovery and Data Mining (KDD)* (2007) pp. 240–249
- 15.59 Z. Duan, B. Pardo: Soundprism: An online system for score-informed source separation of music audio, *IEEE J. Selected Topics Signal Process.* **5**(6), 1205–1215 (2011)
- 15.60 S. Ewert, M. Müller, P. Grosche: High resolution audio synchronization using chroma onset features. In: *Proc. IEEE Int. Conf. Acoustics Speech Signal Process. (ICASSP)* (2009) pp. 1869–1872
- 15.61 C. Joder, S. Essid, G. Richard: A conditional random field framework for robust and scalable audio-to-score matching, *IEEE Trans. Audio Speech Lang. Process.* **19**(8), 2385–2397 (2011)
- 15.62 A. Doucet, N. de Freitas, N.J. Gordon (Eds.): *Sequential Monte Carlo Methods in Practice* (Springer, New York 2001)
- 15.63 M.S. Arulampalam, S. Maskell, N. Gordon, T. Clapp: A tutorial on particle filters for online nonlinear-non-Gaussian Bayesian tracking, *IEEE Trans. Signal Process.* **50**(2), 174–188 (2002)