

JAZZNET: A DATASET OF FUNDAMENTAL PIANO PATTERNS FOR MUSIC AUDIO MACHINE LEARNING RESEARCH

Tosiron Adegbiya

Department of Electrical and Computer Engineering, University of Arizona, Tucson AZ, USA
tosiron@arizona.edu

ABSTRACT

This paper introduces the *jazznet Dataset*, a dataset of fundamental jazz piano music patterns for developing machine learning (ML) algorithms in music information retrieval (MIR). The dataset contains 162520 labeled piano patterns, including chords, arpeggios, scales, and chord progressions with their inversions, resulting in more than 26k hours of audio and a total size of 95GB. The paper explains the dataset’s composition, creation, and generation, and presents an open-source *Pattern Generator* using a method called *Distance-Based Pattern Structures (DBPS)*, which allows researchers to easily generate new piano patterns simply by defining the distances between pitches within the musical patterns. We demonstrate that the dataset can help researchers benchmark new models for challenging MIR tasks, using a convolutional recurrent neural network (CRNN) and a deep convolutional neural network. The dataset and code are available via: <https://github.com/tosiron/jazznet>.

Index Terms— Music information research, machine learning, jazz piano dataset, big data

1. INTRODUCTION

One of the most important and basic needs of machine learning (ML) research in music is the availability of free datasets. Unfortunately, the field of music has lagged behind other fields (like image and speech recognition) in the availability of high-quality datasets. This paper contributes to the growing body of available music audio datasets by presenting the *jazznet dataset*, an extensible dataset of fundamental piano patterns.

There are notable efforts toward creating datasets for MIR, and jazznet is complementary to these efforts. Some existing music datasets include GTZAN [1], MSD [2], AudioSet [3], FMA dataset [4], MusicNet [5], and RWC [6]. The main difference between the jazznet dataset and the majority of existing datasets lies in the approach taken to creating the dataset: an emphasis on the fundamental patterns, rather than complete music pieces. This approach is inspired by how humans effectively learn piano music. Suppose you wanted to learn to play jazz piano; you would learn much faster and more effectively if you first understood the fundamental patterns—*chords*, *arpeggios*, *scales*, and *chord progressions*—on which full music pieces are based [7]. We focus on jazz piano music patterns to make the dataset creation tractable and because of the versatility of jazz music—the patterns in jazz music

encompass several other musical genres (e.g., blues, country, pop, etc.). Given that an important hallmark of jazz music is the variety of expression, the dataset also contains the different forms—or *inversions*—in which the patterns can be played.

The multi-class characteristic of the dataset also aims to model the hierarchy of difficulty humans experience in learning jazz music. For instance, it is easy to differentiate between the *types* of patterns (e.g., chords vs. arpeggios vs. scales). Within each of the classes, it is substantially more difficult to recognize different *modes* of patterns (e.g., major vs. minor or augmented vs. diminished chords). Finally, only a uniquely talented musician with *absolute pitch* [8] can recognize specific modes (e.g., A major vs. C major).

In summary, we make two major contributions. First, we present the jazznet dataset¹, which we have curated based on a wide-ranging survey of jazz education resources and jazz standards. The dataset contains 162520 automatically generated and labeled piano music patterns, which results in 95GB and over 26k hours of audio. The primary objective of the jazznet dataset is to facilitate the development of ML models for challenging MIR tasks. Second, we have developed a *Pattern Generator* that uses *Distance-Based Pattern Structures (DBPS)* to facilitate the generation of piano patterns based on the distance between pitches within the patterns. This approach enables the easy generation of new piano patterns for ML research, and we have open-sourced the Python scripts² that implement it for the convenience of users.

2. THE JAZZNET DATASET

Although complementary to existing music datasets, jazznet is unique in several ways. Structurally, the most related datasets to jazznet are the NSynth [9] and MAPS [10] datasets, with some notable differences. NSynth is created for exploring neural audio synthesis of musical notes with different pitches, timbres, and envelopes. However, it does not contain any musical patterns, only single notes, and although it represents more musical instruments and dynamics than jazznet, it is not comparable in terms of the variety of musical patterns contained in jazznet. On the other hand, the MAPS dataset is frequently used for automatic music transcription [11] and contains isolated

¹<https://doi.org/10.5281/zenodo.7192653>

²<https://github.com/tosiron/jazznet/tree/main/PatternGenerator>

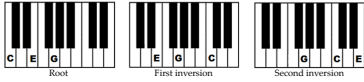


Fig. 1. Illustration of the C major chord and its three forms: root, first, and second inversions

notes, chords, and music pieces. However, it contains fewer chords than jazznet and does not include scales, arpeggios, or chord progressions. Additionally, the MAPS dataset is significantly smaller than jazznet, with 31GB of data compared to 95GB. Most importantly, unlike previous datasets, jazznet is accompanied by piano pattern generators that enable researchers to extend the dataset to include new patterns.

While the dataset is created to be easily extensible, we curated the patterns for the current version by extensively surveying several resources for jazz music performance and education, such as *The Jazz Piano Book* [7], *The Jazz Theory Book* [12]), and more than 50 jazz standards [13, 14]. We also interacted with several jazz musicians to develop a consensus on some of the common fundamental patterns, particularly chord progressions. However, we acknowledge that given the complexity and diversity of jazz music, the dataset does not cover all jazz piano patterns. Nevertheless, the provided pattern generator (Section 2.5) enables users to generate new patterns not in the dataset. In the following sections, we provide a background on music theory necessary for understanding the dataset’s creation, describe the data generation approach, and then present the jazznet dataset, its features, the pattern generator, and potential applications.

2.1. Brief background basic music theory

The dataset is based on the notes of a standard 8-octave (88-key) piano. The main notes of a piano (and most other stringed instruments) are typically represented using 7 letters: *A, B, C, D, E, F, G* (which correspond to the white keys on a piano). The notes in between these notes (i.e., the black keys) are described as *sharp* (\sharp) to the immediately preceding note, or *flat* (*b*) to the immediately succeeding note. For example, the note between C and D is referred to as $C\sharp$ or D_b depending on the context. An *octave* is the interval between one musical pitch and another half or double its frequency. For example, the middle C on the piano (also called $C4$) has a frequency of 261.63 Hz, the next C ($C5$) has a frequency of 523.25 Hz, and the previous C ($C3$) has a frequency of 130.81 Hz. There are 12 ‘half steps’ (also called *semitones*) or 6 ‘full steps’ (also called *tones*) between each note and its octave.

The dataset consists of four types of piano patterns: *chords*, *arpeggios*, *scales* and *chord progressions*. A chord is a harmonic set of pitches or frequencies consisting of multiple notes that are played or heard simultaneously. Chords can be categorized by the number of notes they contain: *dyads* (2-note), *triad* (3-note), *tetrad* (4-note), etc. An arpeggio is a chord in which the notes are played one after another. A scale is a set of notes ordered by the frequency of the notes. The notes in these patterns can be rearranged to create different ‘colors’ of music, called *inversions* (illustrated in Figure 1 using the C major triad). Additionally, one of the first and most important things a

Algorithm 1: Pattern generation via *distance-based pattern structures (DBPS)*

Input: distance = $[d_0, d_1, d_2, \dots, d_n]$; type=“chord” or “arpeggio” or “scale”
Output: Pattern in all keys

```

1 foreach MIDI_pitch in range(24, 109) do
2    $note_0 \leftarrow \text{MIDI\_pitch}$ 
3    $time \leftarrow \text{time} \# \text{current time}$ 
4   for i in range(0, len(distance)) do
5      $note_{i+1} \leftarrow note_i + \text{distance}[i]$ 
6     if type == “arpeggio” or “scale” then
7        $time \leftarrow time + 1$ 
8     end
9   end
10 end

```

new jazz pianist must learn is common chord progressions. Chord progressions are successions of chords that form the foundation of music [12]. They are often represented using Roman numerals, with lowercase letters for minor chords and uppercase letters for major chords. This representation allows progressions to be independent of the specific key that a piece of music is played in.

2.2. Dataset creation via *distance-based pattern structures (DBPS)*

We use a method called *distance-based pattern structures (DBPS)* to automatically generate piano patterns in a flexible and intuitive way. Thanks to the formulaic structure of most piano music patterns, the patterns can be generated based on the distance between the pitches within the patterns. We represent the pitches using MIDI pitch numbers, with a distance of 1 representing the shortest distance between two notes (i.e., a semitone). To illustrate this approach, we provide a simplified pseudocode in Algorithm 1. The algorithm takes the pitch distances from preceding notes as input to generate any pattern. For example, for any major triad with three notes $note_0, note_1, note_2$, the input would be $distance = [4, 3]$. That is, given the base note as $note_0$, $note_1$ is 4 pitches (or 4 semitones) away, and $note_2$ is 3 pitches away. Similarly, for a maj7 tetrad, $distance = [4, 3, 4]$; and so on.

To generate chord progressions, we use a similar approach based on the distances between notes within each chord of the progression. The base note ($note_0$) for each chord is defined by its distance from the reference note (the key in which the progression exists) based on the Roman numeral. For example, the base note of the *ii* chord is two pitches from the reference note, while the base note of the *IV* chord is five pitches away. Once the base note is determined, the chords in the progression can be generated as described in Algorithm 1. For example, the popular jazz chord progression, *ii-V-I*, can be represented by the following pattern structure:

```

chord1: note0=ref+2; note1=note0+3; note2=note1+4
chord2: note0=ref+7; note1=note0+4; note2=note1+3
chord3: note0=ref; note1=note0+4; note2=note1+3

```

The DBPS approach can also be used to generate *altered* or *extended* chords commonly found in jazz music.

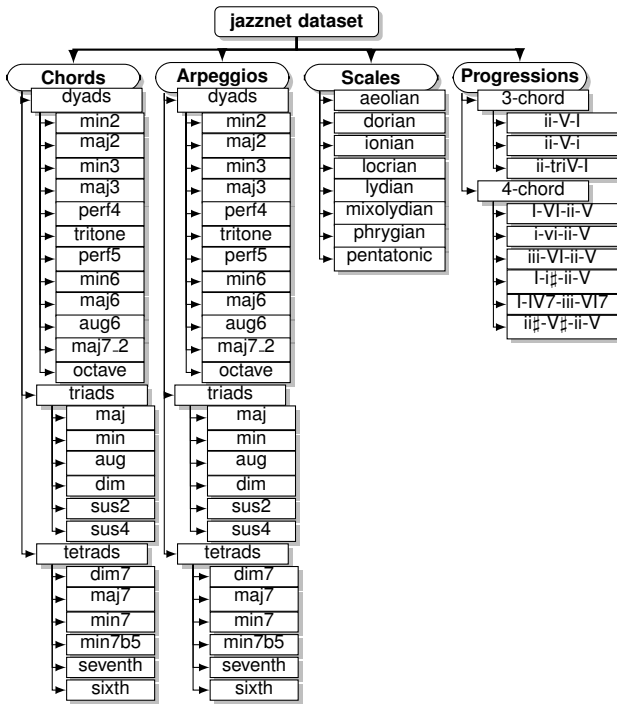


Fig. 2. High-level taxonomy of the jazznet Dataset.

The approach works by imposing the distance required for the alteration or extension to generate the new chord. For example, a dominant 7th extension on a chord with notes $note_0$, $note_1$, and $note_2$, would add a new note $note_3$, three pitches above $note_2$. Adding a major 7th would add a note that is four pitches above $note_2$. Similarly, alterations can be generated by modifying the pitch of a note within the chord. For instance, a b_5 would be created by reducing the pitch of $note_2$ by one.

We created pattern structures for all the piano patterns selected in our survey, and then used Python to develop a pattern generator (Section 2.5) that utilized these structures to generate MIDI format files³ for all keys and inversions. These MIDI files were then manually reviewed for accuracy before being converted to WAV format using the Timidity synthesizer⁴, with a sampling rate of 16KHz and 24-bit depth. To ensure the quality of the generated patterns, we randomly selected WAV files for verification by two jazz pianists, who played the patterns on physical pianos to confirm that the sounds and labels were correct.

2.3. Dataset description and statistics

Figure 2 depicts a high-level taxonomy of the dataset. The dataset contains three kinds of chords and arpeggios: *dyads*, *triads*, and *tetrads*. The scales are dominated by *diatonic* (6-note) scales and include one *pentatonic* (5-note) scale. The dataset is dominated by progressions, of which there are two kinds: 3-chord and 4-chord (tetrad) progressions. All the patterns are present in all their inversions and in all keys of the 88-key piano.

Table 1 summarizes the dataset statistics. The dataset is published in WAV format (MIDI files are also provided) as separate data directories for *types*: chords, arpeggios,

Table 1. Statistics of the jazznet dataset

Type	#modes	#total	time (s)	# hours	size (GB)
Chords	24	5,525	3	259	1
Arpeggios	24	5,525	4; 5; 6	433	1.7
Scales	8	4,590	9; 7	674	6.3
Progressions	9	146,880	7; 10	25,568	85.6
Total	65	162,520		26,934	95

scales, and progressions, and subdirectories for the different *modes*. There are 65 labeled modes detailed as:

- **Chords/Arpeggios** (24 each, appended with “-chord” or “-arpeggio”): 12 *dyads* (1 inversion each): minor 2nd (min2), major 2nd (maj2), minor 3rd (min3), major 3rd (maj3), perfect 4th (perf4), tritone, perfect 5th (perf5), minor 6th (min6), major 6th (maj6), minor 7th or augmented 6th (aug6), major 7th (maj7.2), octave; 6 *triads* (2 inversions each): major (maj), minor (min), augmented (aug), diminished (dim), suspended 2nd (sus2), suspended 4th (sus4); and 6 *tetrads* (3 inversions each): dim7, maj7, min7, min7 b_5 , seventh, sixth.
- **Scales** (8): 5 *diatonics* (6 inversions each): aeolian, dorian, ionian, locrian, lydian, mixolydian, phrygian; 1 *pentatonic* (4 inversions). There are 4590 scales in total.
- **Progressions** (9): 3-chord (64 combinations): ii-V-I, ii-V-i, and ii-triV-I (triV means tritone substitution of V); 4-chord (256 combinations): I-VI-ii-V, i-vi-ii-V, iii-VI-ii-V, I-ii-ii-V, I-IV7-iii-VI7, and ii \sharp -V \sharp -ii-V. There are 146,880 progressions in total.

Each mode subdirectory contains the patterns in each mode labeled according to the specific note/pitch, octave, mode, and inversion. The inversions in the progressions are labeled according to the chord combinations.

The patterns are all recorded at a speed of 60 beats per minute (bpm), with two beats of decay at the end. The chord recordings are all 3 seconds long. The arpeggios and scales are recorded with one note played per beat. The arpeggios range from 4 to 6 seconds long, while the scales range from 7 to 9 seconds long, depending on the number of notes. Progressions are recorded with two chords played per measure and range from 7 to 10 seconds long.

2.4. Suggested subsets

Given the dataset’s imbalance resulting from the size of the chord progressions, we suggest some subsets of the dataset where a more balanced dataset may be required or downloading/using the entire dataset may be impractical (e.g., for initially testing a model). All subsets contain all the chords, scales, and arpeggios (15640 clips). The small, medium, and large subsets add 5876, 14688, and 36720 progressions, for a total of 21516, 30328, and 52360 clips, respectively. All the subsets are pseudo-randomly generated to ensure that all the modes are represented. The meta-data files are in CSV format with a suggested split into train, validation, and test sets using an 80/10/10% split. The validation and test sets are randomly selected from the different modes without overlapping with the training set.

³<https://github.com/MarkCWirt/MIDIUtil>

⁴<http://timidity.sourceforge.net/>

2.5. Pattern generator

We have developed piano *pattern generators* using Python, which allow users to easily create new patterns with ease. Due to space limitations, we omit the full details of the generator functions (the details are on the Github page⁵), but summarize that chords, arpeggios, and scales can be generated by specifying the pitch distances as outlined in Section 2.2. In the same way, chord progressions can be generated by specifying the Roman numerals of the chords and the supported extensions/alterations (e.g., maj7, b5, ♯, etc.). Crucially, the patterns are generated in all keys of an 88-key piano and in all inversions. The generated MIDI files can then be converted into the user’s format of choice, such as WAV, using the user’s choice of tool. The provided tool allows for the creation of numerous new patterns, and the scripts are open-source, allowing for the possibility of expanding the tool to support additional patterns.

2.6. Applications

Jazznet allows the evaluation of models developed for a variety of machine learning MIR tasks. The most straightforward and basic tasks involve **automatic music understanding** based on the dataset’s class hierarchy: type recognition (whether an input is a chord, arpeggio, scale, or chord progression) or mode recognition, e.g., whether an input is augmented (aug) or major (maj) or Ionian or a ii-V-I progression. Specific details of the patterns can also be predicted, like the pitch contents, specific key (e.g., C-maj vs. Eb-maj chord), octaves, and inversions.

Jazznet can be used for more complex MIR tasks. For example, **automatic music transcription** [11, 15] is considered a canonical task in MIR, where the musical components (e.g., chords) in an audio recording are extracted. Most current music transcription models focus on isolated notes or chords present in the music. But additional musical patterns like scales and chord progressions can be predicted using models trained on jazznet. The dataset can also be used for the development of **music recommendation systems** [16, 17]. Recommendation systems can be built based on musical patterns: a listener who likes certain kinds of chords or scales in one song (e.g., a song dominated by the minor scale or minor chords) may similarly like other songs that contain similar chord modes. Related to music understanding, jazznet can be used for **automatic music generation** [18, 19]. A lot of jazz music involves the repetition of jazz chord progressions with scales or arpeggios played over them. For example, the ii-V-I progression, which appears in nearly all jazz standards [7], can be soloed over with each chord’s arpeggios or with matching scales (e.g., the Dorian mode over the ii chord). Jazznet can be used to learn the fundamental patterns in musical pieces and how they occur (e.g., frequency of progressions, scales used with the progressions, etc.).

3. EXPERIMENTS

Experimental setup: We used two simple models to demonstrate the potential of the jazznet dataset for music

⁵<https://github.com/tosiron/jazznet/tree/main/PatternGenerator>

Table 2. Average precision and mAP for mode prediction.

	Chords	Arpeggios	Scales	Progressions	mAP
CRNN	0.28	0.16	0.10	0.67	0.63
M5	0.36	0.09	0.06	0.48	0.30

audio recognition in two tasks, type and mode recognition. The first model is a convolutional recurrent neural network (CRNN), as described in [20]. This model utilizes Mel-spectrograms as input and models long-term dependencies in musical structures using a recurrent neural network (RNN). The second model is the M5 model described in [21]: a deep convolutional neural network that takes time series waveforms as input and requires no preprocessing. Since data preprocessing is often a bottleneck for audio ML, this model provides a sense of the performance achieved using the dataset without any preprocessing. Both models contain 4 convolutional layers with batch normalization and the ReLU activation function. We trained the models with the Adam optimizer and categorical crossentropy loss function, and evaluated their performance using the mean Average Precision (mAP) and AUC (Area Under the Curve) score, following prior work [21], on the medium subset (30328 clips) of the dataset.

Results: Identifying chords, arpeggios, scales, and progressions was a simple task for the models. The CRNN and M5 models had high scores with an average AUC score of 0.984. They consistently recognized all four types, with mAP scores of 0.99 and 0.97. In comparison, two jazz pianists informally attempted to recognize a random sample of the types and both achieved 100% accuracy.

On the other hand, recognizing modes across the 65 modes was more difficult for the models, as summarized in Table 2. The CRNN and M5 models achieved mAP scores of 0.63 and 0.30, respectively, with an average AUC score of 0.04. Both models found the arpeggios and scales to be the most challenging with AP of 0.16 and 0.10 for CRNN and 0.09 and 0.06 for M5. However, the models performed best on the progressions with an AP of 0.67 and 0.48 for CRNN and M5, respectively. The humans also found mode prediction more difficult than type recognition, achieving an accuracy range of 57% to 73%. In summary, the dataset was designed to reflect the hierarchy of difficulty in recognizing musical patterns and provides an opportunity to benchmark models for challenging MIR tasks.

4. CONCLUSION

In this paper, we presented a *jazznet*, a dataset of essential jazz piano music patterns and an open-source pattern generator, enabling researchers to benchmark machine learning models for complex music information retrieval (MIR) tasks. Our aim is for this dataset to contribute to advancing machine learning research in MIR. In future work, we plan to expand the dataset by including more musical auditory attributes such as dynamics and rhythmic variations, as well as more complex patterns like 5-note chords and longer chord progressions. Additionally, we intend to explore more sophisticated models to improve performance on tasks using the dataset and also further investigate the potential of the DBPS approach for automatically generating different kinds of data.

5. REFERENCES

- [1] George Tzanetakis and Perry Cook, “Musical genre classification of audio signals,” *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293–302, 2002.
- [2] Thierry Bertin-Mahieux, Daniel PW Ellis, Brian Whitman, and Paul Lamere, “The million song dataset,” 2011.
- [3] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio Set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [4] Michaël Defferrard, Kirell Benzi, Pierre Vandergheynst, and Xavier Bresson, “FMA: A dataset for music analysis,” *arXiv preprint arXiv:1612.01840*, 2016.
- [5] John Thickstun, Zaid Harchaoui, and Sham Kakade, “Learning features of music from scratch,” *arXiv preprint arXiv:1611.09827*, 2016.
- [6] Masataka Goto, Hiroki Hashiguchi, Takuichi Nishimura, and Ryuichi Oka, “RWC music database: Music genre database and musical instrument sound database,” 2003.
- [7] Mark Levine, *The Jazz Piano Book*, “O’Reilly Media, Inc.”, 2011.
- [8] Diana Deutsch, “The enigma of absolute pitch,” *Acoustics Today*, vol. 2, no. 4, pp. 11, 2006.
- [9] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan, “Neural audio synthesis of musical notes with wavenet autoencoders,” in *International Conference on Machine Learning*. PMLR, 2017, pp. 1068–1077.
- [10] Valentin Emiya, Nancy Bertin, Bertrand David, and Roland Badeau, “Maps-a piano database for multipitch estimation and automatic transcription of music,” *Research Report, Inria-00544155*, 2010.
- [11] Emmanouil Benetos, Simon Dixon, Zhiyao Duan, and Sebastian Ewert, “Automatic music transcription: An overview,” *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20–30, 2018.
- [12] Mark Levine, *The Jazz Theory Book*, “O’Reilly Media, Inc.”, 2011.
- [13] Ted Gioia, *The jazz standards: A guide to the repertoire*, Oxford University Press, 2021.
- [14] Hal Leonard Corp, *First 50 jazz standards you should play on piano*, Hal Leonard, 2017.
- [15] Emmanouil Benetos, Simon Dixon, Dimitrios Gianoulis, Holger Kirchhoff, and Anssi Klapuri, “Automatic music transcription: challenges and future directions,” *Journal of Intelligent Information Systems*, vol. 41, no. 3, pp. 407–434, 2013.
- [16] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminskas, “Music recommender systems,” in *Recommender Systems Handbook*, pp. 453–492. Springer, 2015.
- [17] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen, “Deep content-based music recommendation,” *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [18] Dorien Herremans, Ching-Hua Chuan, and Elaine Chew, “A functional taxonomy of music generation systems,” *ACM Computing Surveys (CSUR)*, vol. 50, no. 5, pp. 1–30, 2017.
- [19] Filippo Carnovalini and Antonio Rodà, “A multilayered approach to automatic music generation and expressive performance,” in *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*. IEEE, 2019, pp. 41–48.
- [20] Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho, “Convolutional recurrent neural networks for music classification,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 2392–2396.
- [21] Wei Dai, Chia Dai, Shuhui Qu, Juncheng Li, and Samarjit Das, “Very deep convolutional neural networks for raw waveforms,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 421–425.