MIREX 2026 Overview

ISMIR 2026 Late-Breaking/Demo · Abu Dhabi

Future MIREX Team

All 2026 results on the MIREX wiki

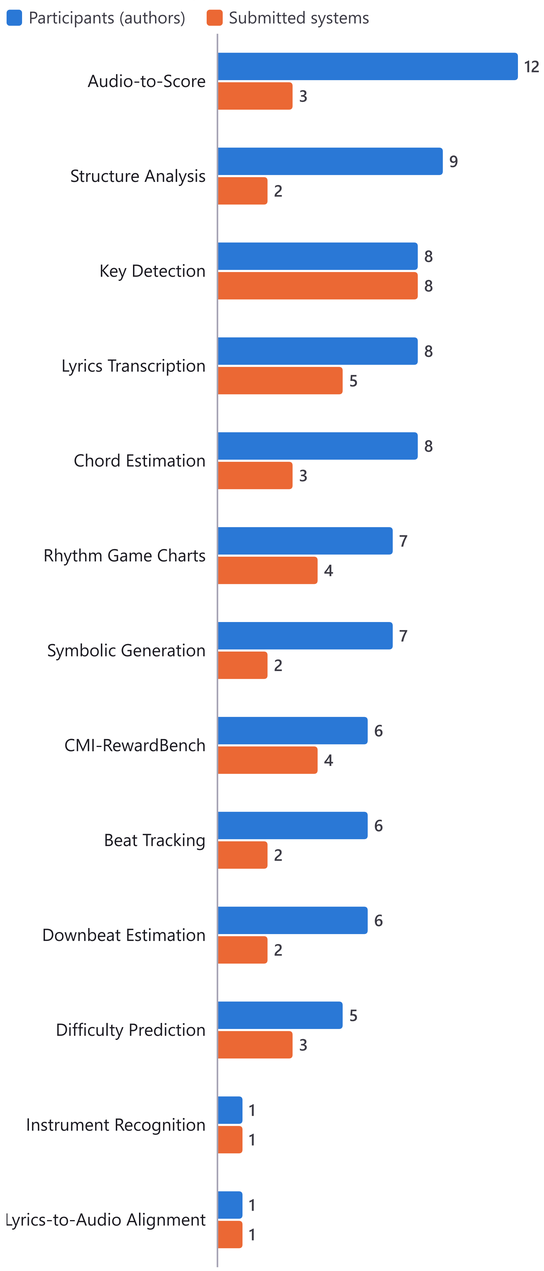
MIREX 2026 at a glance

13 tasks

40 submitted systems

63 authors

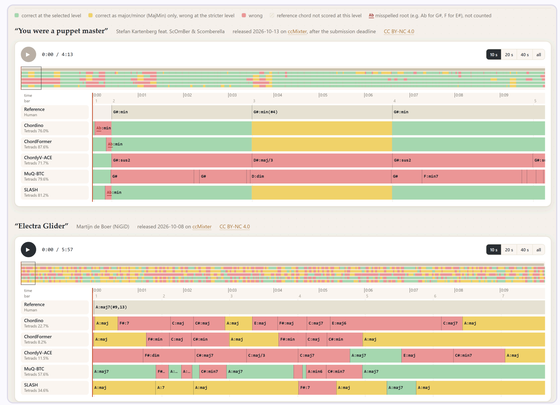
Participants and systems per task



Baselines are not counted. Task descriptions on this page are quoted from the 2026 task pages of the MIREX wiki.

Audio Chord Estimation

This task requires participants to extract or transcribe a sequence of chords from an audio music recording. This includes the full characterisation of chords – root, quality, and bass note – as well as their chronological order, including specific onset times and durations.



Interactive demo
Two Creative Commons songs released after the submission deadline, with every system's chords.
futuremirex.com/demos/audio-chord-estimation-2026/

Audio Key Detection

Audio Key Detection aims to identify the musical key (e.g., C major, A minor) of an audio recording. This involves determining both the tonic (root pitch) and the mode (major or minor) from the audio signal.

Audio Beat Tracking

The aim of the automatic beat tracking task is to track each beat location in a collection of sound files. The algorithms are evaluated in terms of their accuracy in predicting beat locations annotated by a group of listeners.

Audio Downbeat Estimation

The aim of the automatic downbeat estimation task is to identify the locations of downbeats in a collection of sound files. While this is similar to the Audio Beat Tracking task, here the aim is to find the first beat of each bar (measure) rather than all beat times.

Music Audio Generation

The MIREX 2026 **Cover Song Generation** task asks an AI music system to remake a full song according to a style or arrangement prompt. The output can change the style, instruments, arrangement, timbre, or tempo, but it should still be recognizable as a version of the source song.

Music Structure Analysis

The aim of the MIREX Music Structure Analysis task is to identify and label key structural sections in musical audio. For MIREX 2026, participants are required to segment musical audio and classify each segment into one of seven functional categories: 'intro', 'verse', 'chorus', 'bridge', 'inst' (instrumental), 'outro', or 'other'.

Lyrics-to-Audio Alignment

The task of automatic lyrics-to-audio alignment has as an end goal the synchronization between an audio recording of singing and its corresponding written lyrics. The algorithm receives mixed singing audio (singing voice + musical accompaniment) and its lyrics at word level, and outputs the onset and offset timestamps (seconds) of each word.

```
Lyrics:  no more carefree ...
Output:  0.123  0.798  no
         0.798  1.123  more
         1.345  2.176  carefree
```

Lyrics Transcription

The task of Lyrics Transcription aims to identify the words from sung utterances, in the same way as in automatic speech recognition. For this year's edition, the input is always a polyphonic mix (singing voice + musical accompaniment).

Audio Instrument Recognition

The task is clip-level multi-label instrument recognition. Given a music audio excerpt, a submitted system should predict which instruments are present in the excerpt.

Audio-to-Score Transcription

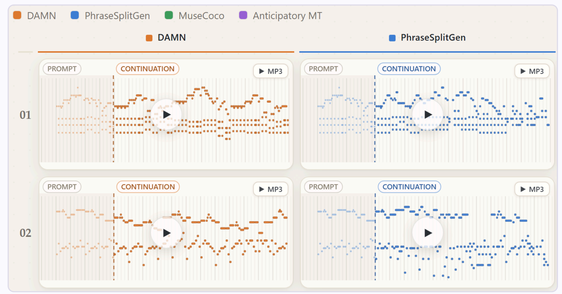
Participating systems receive an audio recording of a polyphonic music piece and must output a digital symbolic score that can be read by a musician or notation program: a kern score encoding pitches, durations, accidentals, staves, meter, clef, key signature, and bar structure. Two tracks: **Staves-Informed A2S** (staves metadata is provided) and **Blind A2S** (only the audio recording is provided).

Music Performance Difficulty Prediction

Participants submit systems that take a solo piano audio recording as input and output a predicted difficulty score. This shared task establishes the first community-wide benchmark for piano difficulty estimation from audio.

Symbolic Music Generation

For this year's challenge, the selected task is **Piano Music Continuation**. Given a 4-measure piano prompt (plus an optional pickup measure), the goal is to generate a 12-measure continuation that is musically coherent with the prompt, forming a complete 16-measure piece.



Interactive demo
8 prompts × 4 systems, playable in the browser (excerpt shown).
futuremirex.com/demos/symbolic-music-generation-2026/

Rhythm Game Chart Generation

A rhythm game chart consists of a sequence of timed note events synchronized with an audio track, defining what actions players should perform and when. This year's task is generating **osu!taiko** beatmaps: osu!taiko is a game mode in osu! inspired by Taiko no Tatsujin, which features a horizontal conveyor belt of drum notes moving towards a judgment circle.



Hit circles | Drumroll | Spinner

Images: MIREX wiki task page.

Music Evaluation via CMI-RewardBench

Participants are invited to submit evaluation scripts or models that output preference scores given a prompt and a pair of generated music tracks. The benchmark evaluates the systems on their ability to mirror human judgment.

```
Prompt:  "ambient serene instrumental"
Input:   Candidate A, Candidate B (generated tracks)
Output:  {"battle_uuid": "toy_example_001",
          "preferred_candidate": "A"}
```