

Control Improvisation With Application to Music

Alexandre Donzé¹, Sophie Libkind², Sanjit Seshia¹, and David Wessel¹

University of California, Berkeley

Swarthmore College

September 29, 2013

Control Improvisation Problem

Generic Goal:

- ▶ Generate control inputs in a changing environment
- ▶ “improvised” \implies randomized
- ▶ Soft constraint: stay “close” to some nominal controller
- ▶ Hard constraints: always satisfy some specifications

In the musical context:

- ▶ Generate random melodies in a given harmonic context
- ▶ Generated melodies that “sound similar” to a reference melody

Control Improvisation Problem

Generic Goal:

- ▶ Generate control inputs in a changing environment
- ▶ “improvised” \implies randomized
- ▶ Soft constraint: stay “close” to some nominal controller
- ▶ Hard constraints: always satisfy some specifications

In the musical context:

- ▶ Generate random melodies in a given harmonic context
- ▶ Generated melodies that “sound similar” to a reference melody

Control Improvisation Problem

Formally, given

- ▶ A specification FSM \mathcal{A}^s ,
- ▶ A reference accepting word w_{ref} ,
- ▶ A similarity (or “creativity”) measure d between words,
- ▶ A target distance (or “creativity”) d_r from w_{ref}
- ▶ Any $\epsilon, \delta > 0$

Find a FSA \mathcal{A}^c such that

$$w \in \mathcal{L}(\mathcal{A}^s || \mathcal{A}^c) \implies Pr(|d(w, w_{\text{ref}}) - d_r| < \delta) > 1 - \epsilon$$

The target creativity controls how far the improvisation should be from w_{ref}

Solving the Control Improvisation Problem

General Idea


- ▶ Use learning to create a generator automaton \mathcal{A}_g from w_{ref}
- ▶ Use supervisory control \mathcal{A}_{sc} of the generator to enforce the specification, so that $\mathcal{A}_c = \mathcal{A}_g || \mathcal{A}_{sc}$
- ▶ Tune probabilistic transitions to meet the creativity criterion

Factor Oracle-based improvisations

1. **Input:** a training melody w , represented as a finite sequence of notes, e.g., `abbbbaab`
2. **Learning:** construct a compact structure (factor oracle) representing sub-sequences of w , e.g.,




3. **Generation:** assign probabilities to transitions of the factor oracle and loop on repeated suffixes



Factor Oracle-based improvisations

1. **Input:** a training melody w , represented as a finite sequence of notes, e.g., abbbaab
2. **Learning:** construct a compact structure (factor oracle) representing sub-sequences of w , e.g.,




3. **Generation:** assign probabilities to transitions of the factor oracle and loop on repeated suffixes




Factor Oracle-based improvisations

1. **Input:** a training melody w , represented as a finite sequence of notes, e.g., `abbbaab`
2. **Learning:** construct a compact structure (factor oracle) representing sub-sequences of w , e.g.,



3. **Generation:** assign probabilities to transitions of the factor oracle and loop on repeated suffixes



Properties of Factor Oracles [Assayag and Dubnov, 2004]

A FO reproduces factors of the original string with high probabilities.

With low probabilities, it branches away, generating creativity.

Advantage

Unsupervised learning: no structure is assumed on the input

Disadvantage

Branching likely breaks the underlying structure of the original melody and creates violations of musical specifications

Musical Specifications

The image displays two staves of musical notation for the piece "It Don't Mean a Thing" by Duke Ellington. The top staff is in G minor (one flat) and 4/4 time, with a key signature of one flat. It features four measures with the following chords: Gm, Eb7, D7+, and Gm7. The bottom staff is in C major (no sharps or flats) and 4/4 time, with a key signature of one flat. It features four measures with the following chords: C7, F7, Bb6/B, and D7+. The notation includes various note values, rests, and slurs, illustrating the complex rhythmic and pitch specifications of the piece.

It Don't Mean a Thing by Duke Ellington

- ▶ **Rhythm specifications:** Jazz melodies can be decomposed into licks separated by rests. Licks can be forced to begin on certain beats.
- ▶ **Pitch specifications:** Notes that sound "weird" with a chord cannot begin a lick and must approach a "good" sounding note.

In practice, we decouple generation of note durations (rhythm) and pitches

Musical Specifications

The image displays two staves of musical notation for the piece "It Don't Mean a Thing by Duke Ellington". The music is in 4/4 time and the key signature has two flats (B-flat and E-flat). The first staff contains the following notes and chords: Gm (G-flat major), Eb7 (E-flat dominant 7), D7+ (D dominant 7 with a sharp 9), and Gm7 (G-flat major 7). The second staff contains the following notes and chords: C7 (C dominant 7), F7 (F dominant 7), Bb6/B (B-flat major 6 over B), and D7+ (D dominant 7 with a sharp 9). The notes are primarily eighth and quarter notes, with some beamed eighth notes and a few rests.

It Don't Mean a Thing by Duke Ellington

- ▶ **Rhythm specifications:** Jazz melodies can be decomposed into **licks** separated by rests. Licks can be forced to begin on certain beats.
- ▶ **Pitch specifications:** Notes that sound "weird" with a chord cannot begin a lick and must approach a "good" sounding note.

In practice, we decouple generation of note durations (rhythm) and pitches

Musical Specifications

The image shows two staves of musical notation for the jazz standard 'It Don't Mean a Thing' by Duke Ellington. The key signature is B-flat major (two flats) and the time signature is 4/4. The first staff contains the first four measures of the melody, with chords Gm, Eb7, D7+, and Gm7 indicated above. The second staff contains the next four measures, with chords C7, F7, Bb6/B, and D7+ indicated above. The notation includes various note values, rests, and slurs.

It Don't Mean a Thing by Duke Ellington

- ▶ **Rhythm specifications:** Jazz melodies can be decomposed into **licks** separated by rests. Licks can be forced to begin on certain beats.
- ▶ **Pitch specifications:** Notes that sound “weird” with a chord cannot begin a lick and must **approach** a “good” sounding note.

In practice, we decouple generation of note durations (rhythm) and pitches

Musical Specifications

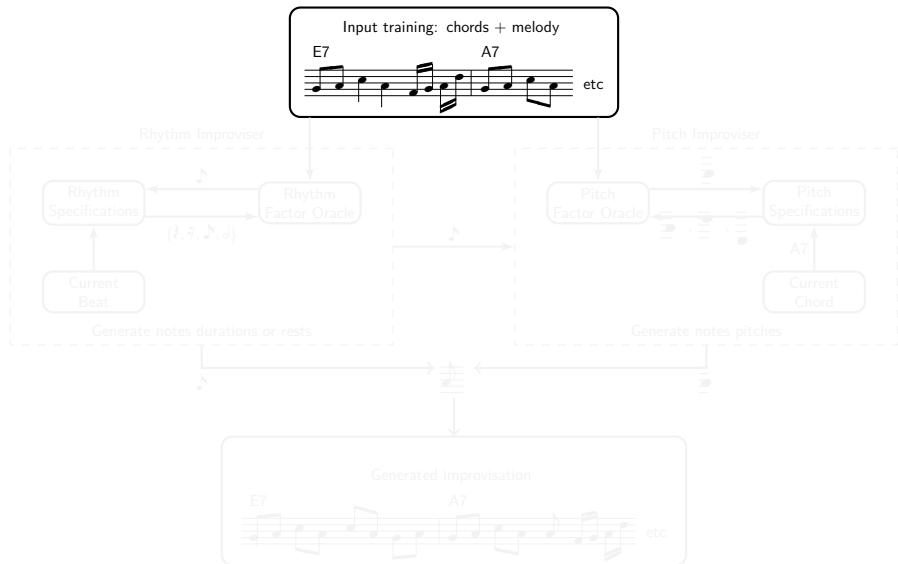
The image displays two staves of musical notation for the piece "It Don't Mean a Thing" by Duke Ellington. The key signature is B-flat major (two flats) and the time signature is 4/4. The first staff contains the following chords: Gm, Eb7, D7+, and Gm7. The second staff contains the following chords: C7, F7, Bb6/B, and D7+. The notation includes various note values, rests, and articulation marks such as slurs and accents.

It Don't Mean a Thing by Duke Ellington

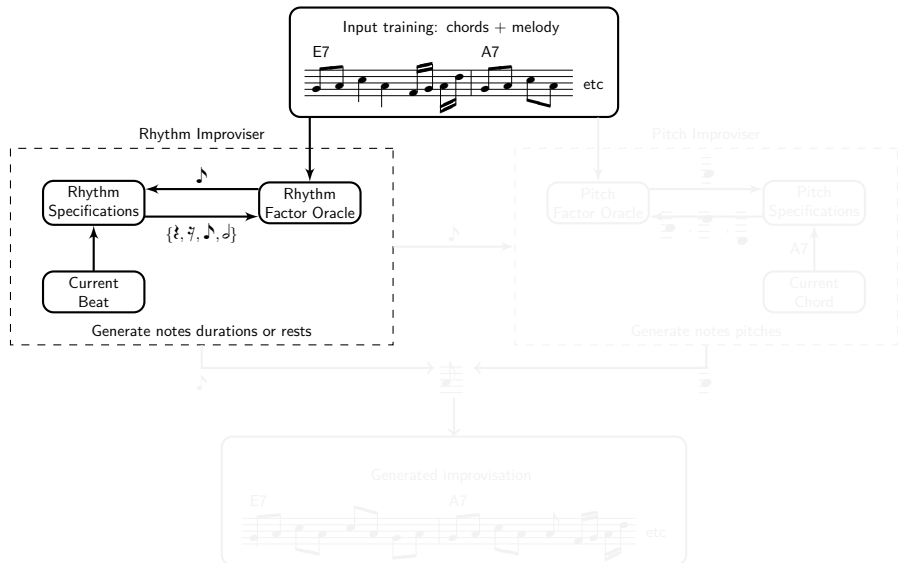
- ▶ **Rhythm specifications:** Jazz melodies can be decomposed into **licks** separated by rests. Licks can be forced to begin on certain beats.
- ▶ **Pitch specifications:** Notes that sound “weird” with a chord cannot begin a lick and must **approach** a “good” sounding note.

In practice, we decouple generation of note durations (rhythm) and pitches

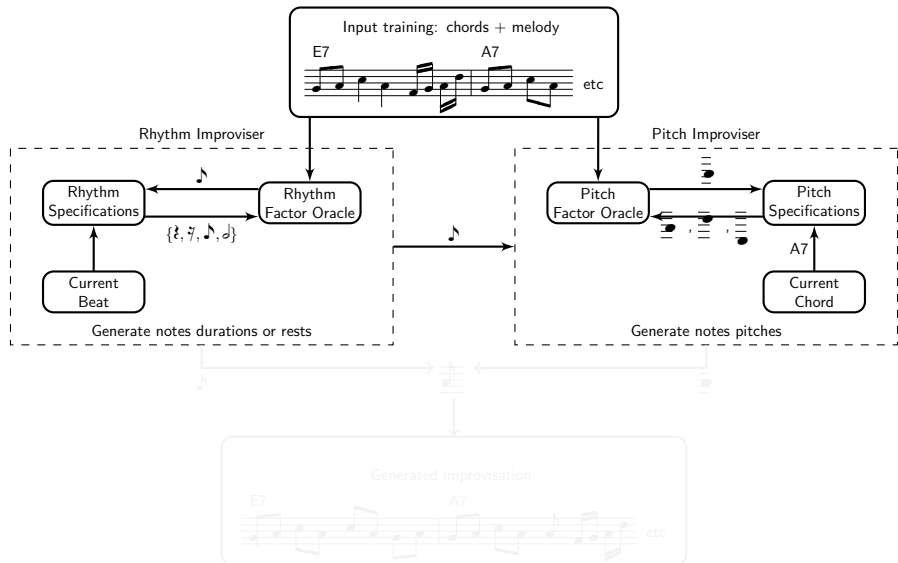
Improviser



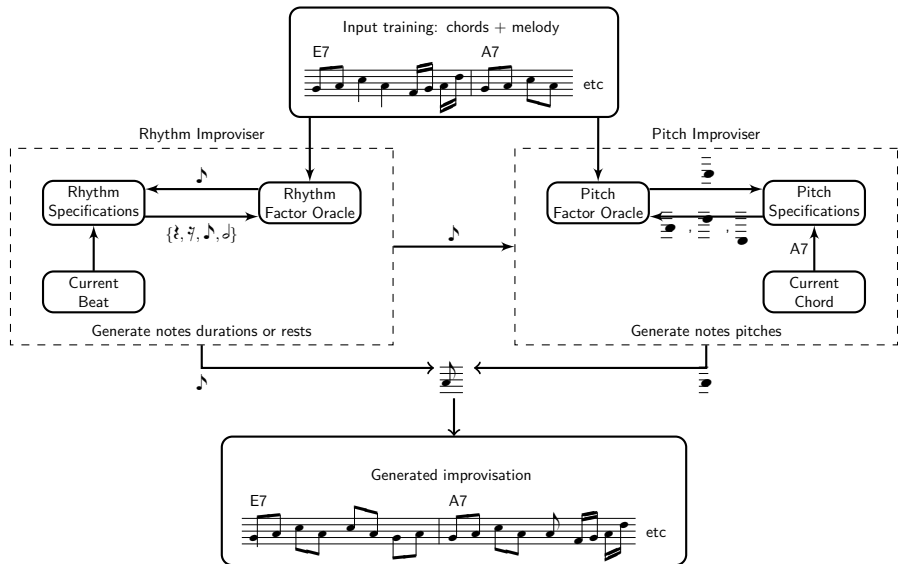
Improviser



Improviser



Improviser



Measuring Creativity

Two sources of divergence:

- ▶ Replication probability (probability of not branching in the FO)
- ▶ To avoid deadlocks (no note proposed by the generator satisfy specs), we allow to pick notes not proposed by the FO

To measure divergence, we use distances based on compression algorithms.

Compression and Complexity

$K(y)$ is the **Kolmogorov complexity** of an object y is the length of the shortest compressed code to which it can be losslessly reduced.

We define a measure of creativity which is a variant of the Normalized Compression Distance:

Creativity Divergence (CD)

Approximates amount of information in y not in x .

$$CD(x, y) = \frac{K(y|x)}{K(y)}$$

$K(y) \approx C(y)$ where

$C(y) = \text{length of } \textit{compress}(y)$

for a compression algorithm *compress*.

Compression and Complexity

$K(y)$ is the **Kolmogorov complexity** of an object y is the length of the shortest compressed code to which it can be losslessly reduced.

We define a measure of creativity which is a variant of the Normalized Compression Distance:

Creativity Divergence (CD)

Approximates amount of information in y not in x .

$$CD(x, y) = \frac{K(y|x)}{K(y)}.$$

$K(y) \approx C(y)$ where

$C(y) = \text{length of } \textit{compress}(y)$

for a compression algorithm *compress*.

Compression and Complexity

$K(y)$ is the **Kolmogorov complexity** of an object y is the length of the shortest compressed code to which it can be losslessly reduced.

We define a measure of creativity which is a variant of the Normalized Compression Distance:

Creativity Divergence (CD)

Approximates amount of information in y not in x .


$$CD(x, y) = \frac{K(y|x)}{K(y)}.$$

$K(y) \approx C(y)$ where

$$C(y) = \text{length of } \textit{compress}(y)$$

for a compression algorithm *compress*.

Evaluating Improvisations



Average creativity with respect to rhythm of 100 improvisations generated by supervise factor oracle.

Controlling Similarity via Compression-based Improvisation

Why?

- ▶ Dictionary based compression relies on pattern finding/matching.
- ▶ We can control creativity by controlling the difference between $C(improv)$ and $C(orig + improv)$.

Idea

Given an original training melody $orig$ and its compressed form we can synthesize an improvisation $improv$ such that

- ▶ $improv$ meets specifications
- ▶ $CD(orig, improv)$ is in a specified range

Controlling Similarity via Compression-based Improvisation

Why?

- ▶ Dictionary based compression relies on pattern finding/matching.
- ▶ We can control creativity by controlling the difference between $C(improv)$ and $C(orig + improv)$.

Idea

Given an original training melody $orig$ and its compressed form we can synthesize an improvisation $improv$ such that

- ▶ $improv$ meets specifications
- ▶ $CD(orig, improv)$ is in a specified range

Conclusion and Future Work

Summary

- ▶ We implemented a supervised musical improviser based on Factor oracles
- ▶ We proposed a divergence “creativity” measure
- ▶ We derived a new improvisation scheme controlling creativity divergence

Future work

- ▶ Further experiments with listeners exposed to our creativity measure
- ▶ Real-time (current work with Max/MSP and Ptolemy via OSC)
- ▶ Application of control improvisation to other domains

Conclusion and Future Work

Summary

- ▶ We implemented a supervised musical improviser based on Factor oracles
- ▶ We proposed a divergence “creativity” measure
- ▶ We derived a new improvisation scheme controlling creativity divergence

Future work

- ▶ Further experiments with listeners exposed to our creativity measure
- ▶ Real-time (current work with Max/MSP and Ptolemy via OSC)
- ▶ Application of control improvisation to other domains