CopyForward: Point-Set Matching for Predicting Patterns

Timothy de Reuse
timothy.dereuse@mcgill.ca

Ichiro Fujinaga
ichiro.fujinaga@mcgill.ca

McGill — Schulich School of Music / École de musique Schulich
CIRMMT — Centre for Interdisciplinary Research in Music Media and Technology
DDMAL — Distributed Digital Music Archives & Libraries Lab

In popular music, future events can often be predicted by copying past events and translating them forward in time.

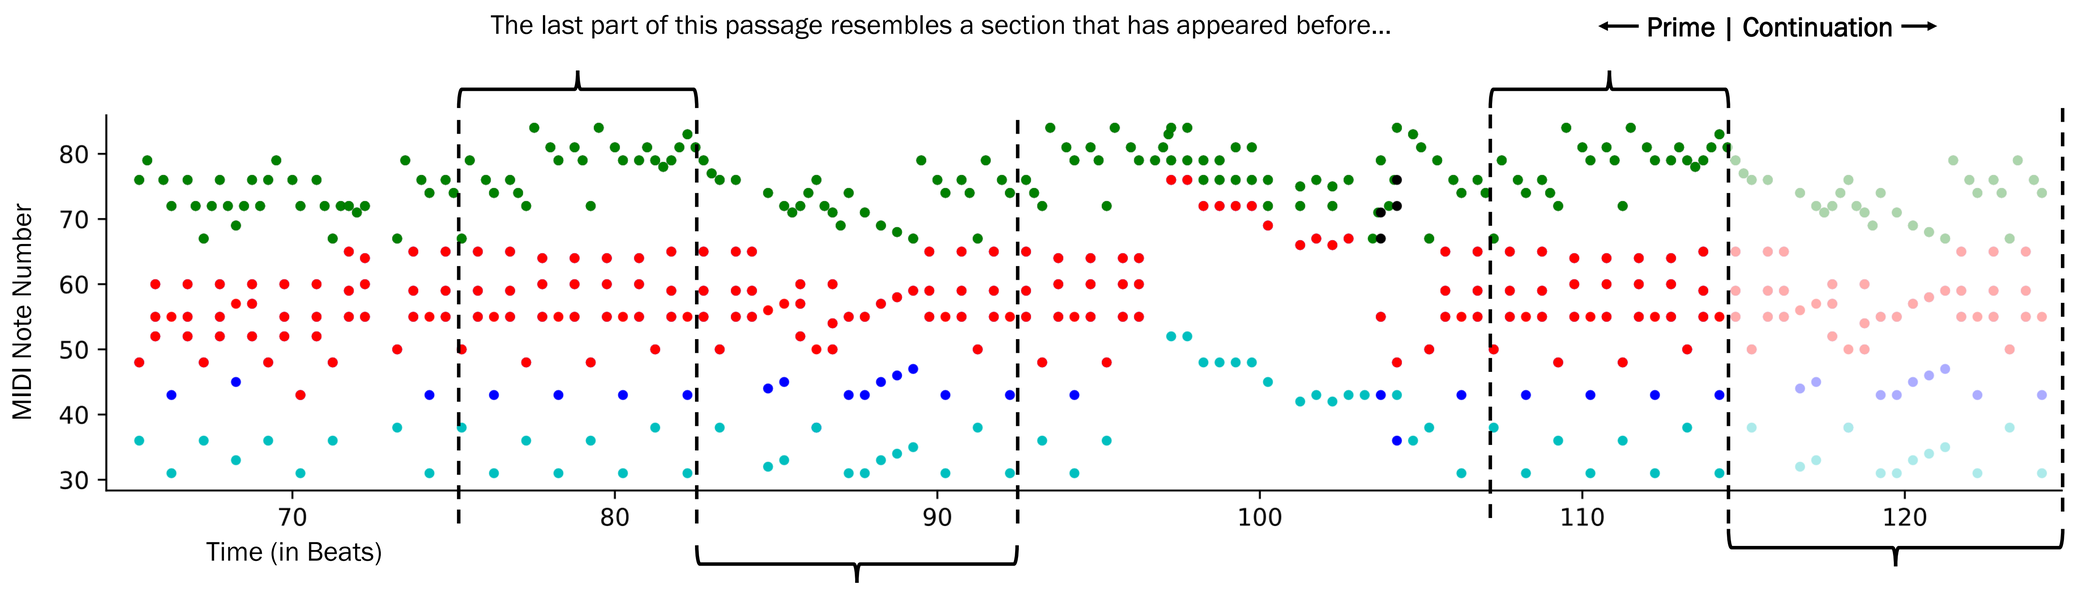

INTRO

- Task: given a *prime* of symbolic music, predict a *continuation*
- Our assumption: the continuation is likely to be **a repetition of material in the prime**
- Task reduces to finding **a single translation vector** that defines what material to copy to the end of the prime

METHOD

- Look at the last few beats of the prime: **does that material occur previously**?
- Look through prime for the **best match to those last few beats** and take whatever occurs after it
- Exceptions for degenerate cases (not enough notes in prime, no good match, desired continuation too long)

RESULTS

- Tested on the Patterns for Prediction Development Dataset (PPDD)
- F1 scores for generated pitch-onset time pairs on 1000 randomly selected entries:
  - Monophonic: **0.493**
  - Polyphonic: **0.445**
- This is only slightly better than the order-1 Markov Models for the monophonic case, but significantly better for the polyphonic case**, especially for longer continuations**

DISCUSSION

- More complicated methods of choosing what part of the prime to copy **failed to yield consistently better results**
- Predictions are usually either completely right or completely wrong
- We hypothesize **an upper limit** to how well a method like this can perform without incorporating more sophisticated musical models

ALGORITHM

1) For:
   - Input, in (MIDI note number, MIDI channel, onset time) triples, lying within time interval $[p_{start}, p_{end}]$
   - Desired continuation length c,
   - Window length $w < c$,
2) Split the prime into 2 parts:
   - *Fixed window*: all events in prime lying in the time interval $[p_{end} - w, p_{end}]$
   - *Sliding window*: all events in prime lying in the time interval $[p_{start}, p_{end} - w)$
3) Find the vector $\boldsymbol{v}$ such that translating the sliding window by $\boldsymbol{v}$ maximizes the number of coinciding points between the two windows
4) Extract all points lying in the time interval $[p_{end}, p_{end} + c]$ from the translated sliding window; this is the predicted continuation

- Through brute-force testing on the PPDD, the best value for w was found to be 8 quarter-note beats.

	Polyphonic			Monophonic		
Algorithm	Recall	Prec.	F1	Recall	Prec.	F1
CopyForward	0.451	0.463	0.445	0.496	0.503	0.493
"Cheating"	**0.562**	**0.553**	**0.547**	**0.645**	**0.644**	**0.638**

We compare this method to one that can "cheat" by always choosing the translation vector $\boldsymbol{v}$ that would give it the best predicted continuation.

Music Information Retrieval Evaluation eXchange (MIREX 2019)
20th International Society for Music Information Retrieval Conference
Delft, Netherlands, November 4-8, 2019

Code and writeup available at
github.com/timothydereuse/copy-forward