

# Math 412 Fall 2014 Data Expeditions Report: Music + Topology

Chris Tralie

April 5, 2015

## 1 Conceptual Overview

This data expedition introduced students to “sliding windows and persistence” on time series data ([1]), which is an algorithm to turn one dimensional time series into a geometric curve in high dimensions, and to quantitatively analyze hybrid geometric/topological properties of the resulting curve such as “loopiness” and “wiggleness.”

In this expedition, musical audio data was the 1D time series in question, which is a fun and relatable way to explore these complicated time series analysis algorithms. Musical audio data is high dimensional (44100 samples per second), noisy, complex, and highly repetitive. It’s the repetitive nature and the global/local descriptions of music that have been of particular interest in my personal research on music information retrieval. As a result, my data expedition was more about exploring “complex data” as opposed to “big data.” Therefore, most of my efforts curating the data were directed at creating a web-based graphical user interface, which I called “Loop Ditty” (<http://www.loopditty.net>), which turned music into visual curves in a way that would allow students to visually explore and manipulate complex patterns in the sound that would be difficult to glean from the waveform alone. Figure 1 shows a screenshot from the LoopDitty software

In addition to visual exploration in 3D, students also used tools from topological data analysis to quantitatively analyze the musical curves in high dimensions. Their task was to make observations in both domains.

## 2 Student Assignment Specifics

### 2.1 Assignment Summary

Students experimented mainly with one parameter, sliding window size, on 30 second excerpts from a rock song and a classical music song. One of the observations students made was that for a smaller window size, the curves had a higher velocity and also tended to concentrate more in one region, while for a larger window size, curves moved more slowly and tended to span more distinct regions of space. Figure 2 shows this phenomenon on Ozzy Osbourne’s “Crazy Train.” Students also used some technology from topological data analysis called a “one dimensional Rips

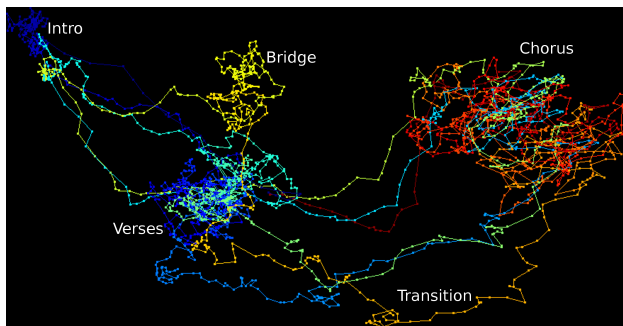


Figure 1: An example of running LoopDitty on Michael Jackson’s “Bad.” Cool colors indicate regions towards the beginning of the song, and hot colors indicate regions towards the end. The verse and chorus regions cluster together, and a bridge (transition) can be seen as a loop that goes away from the main theme and then returns

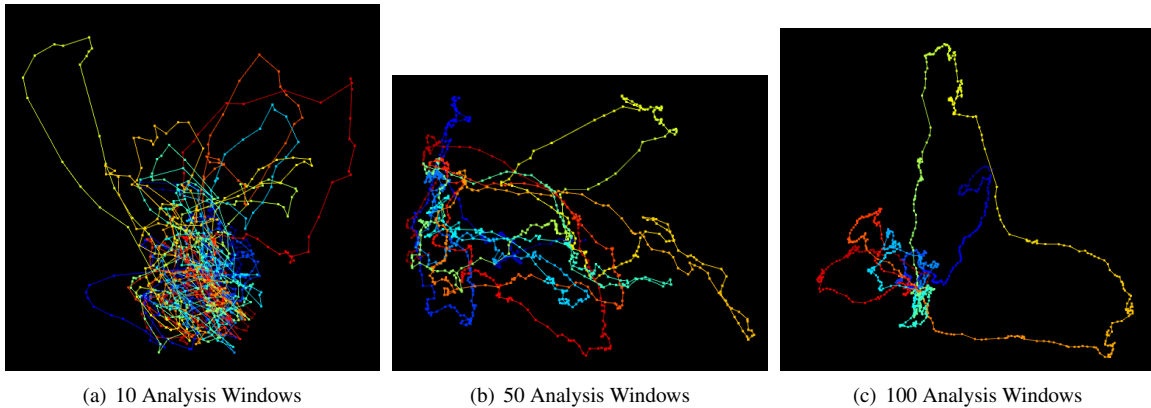


Figure 2: Screenshots from LoopDitty: Varying the sliding window size for the embedding

filtration” to quantitatively describe the size of loops in the curves. With the help of LoopDitty, they observed that the loops in “Crazy Train” corresponded to places in the song where the vocals came in.

Students also used something called a “0D Morse Filtration” to quantify how much the curves wiggled locally in time, and they observed that classical music wiggles less than rock music, both by looking at the Morse filtrations and by looking in LoopDitty. This allowed them to see that they could quantify a qualitative observation about the relative smoothness of classical music with respect to rock music in quantitative geometric terms.

## 2.2 Assignment Results/Observations

The median score on the assignment was a 9/9. I awarded extra credit to students who went onto <http://www.loopditty.net> and tried to make observations about geometric patterns in an entire song of their choice. The students’ observations are still there in the comments section of that web site. I also awarded extra credit to students who went above and beyond with their reports and showed many screenshots from LoopDitty.

One of the themes in this expedition was that visualization is a useful tool, but we can only visualize in three dimensions, while much of the data we care about lives in very high dimensions. Thus, with projections of high dimensional curves down to 3D, some high dimensional detail is necessarily lost, so it is important to be able to use the tools from topological data analysis to say something that perhaps would be missing in the visual representation. Some very perceptive students pointed out that it was slightly unfair for me to ask them to talk about the velocity of curves by just looking at Loop Ditty, as the projection down to 3 dimensions may have collapsed some of the components along which the curves moved the fastest. I admitted to the class that, indeed, this question was not completely fair for that reason, and I used it as an opportunity to highlight to the entire class the potential pitfall of relying too much on projected visualizations (LoopDitty in this case).

Another theme was more specific to the method of turning a 1D time series into a point cloud in high dimensions. There was one question that only three students got correct, which involved a subtle observation about loop formation, but which basically boiled down to an observation that the curve moved faster when the vocals came in. This was an opportunity for me to really test the students’ understanding about the embedding, and to show them how notions of time are actually lost in this representation and they have to be explicitly added back.

Finally, one of the students noticed something in LoopDitty which I had not previously, which ended up informing my research. As I mentioned before, in the “Crazy Train” example, vocals tended to come out in loops. This particular student noticed that the loops for rhyming vocals tended to be closer to each other. This actually makes sense, because the “Mel-Frequency Cepstral Coefficient” features I used in part of the preprocessing were designed to pick up on speech patterns.

Overall, this assignment was clearly a fruitful two-way interaction between me and the students. I believe I got them excited about approaching data both visually and geometrically, and I hope some of them will keep these analysis techniques in the back of their mind and use them someday in their own projects where they are applicable. At the same time, they got me to reflect more deeply on some of the trade-offs between using visual representations and describing data with more abstract high dimensional means, which has fed back into my research.

### 3 Further Details

The full assignment, which includes more technical details on the process, can be found at [http://www.ctralie.com/Teaching/Math412\\_F2014\\_MusicAssignment](http://www.ctralie.com/Teaching/Math412_F2014_MusicAssignment).

Note that I also created a complementary assignment for an electrical engineering audience (ECE 381: Digital Signal Processing) entitled “Musical Pitches and Chroma Features,” which was designed to teach students a robust signal processing approach for extracting musical notes from a 1D time series, but due to schedule conflicts this assignment never materialized. However, I did still create a web page for the assignment, which can be found at [http://www.ctralie.com/Teaching/ECE381\\_DataExpeditions\\_Lab1/](http://www.ctralie.com/Teaching/ECE381_DataExpeditions_Lab1/)

### References

- [1] Jose A Perea and John Harer. Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, pages 1–40, 2013.