

[SQUEAKING]

[CLICKING]

[RUSTLING]

[CLICKING]

MICHAEL SCOTT ASATO CUTHBERT: Hello, computational music researchers. We're going to continue our study of corpora and what we can do with the corpus of music by talking about some research that I did with Sophia H. Sun, who was a visiting student at MIT in 2016-2017 and graduated in 2018 from Wellesley College. We've published our work as a motion painting, lyric, affect, and musical relationships in a large lead sheet corpus.

So this is work that's about looking at a large musical corpus and also looking at words and how music and words interact. Words and musical moments are generally both felt to have intrinsic emotions, but while the emotional context of text is determined at least in part from the semantics that is the meaning of individual words and sentences, the emotional content of individual musical parameters is far less well defined. This lack of definition is especially acute in written musical scores that are separated from the aural and visual components of particular performances or recordings.

So a major question for music perception or the psychology of music is do particular musical features such as consonance and dissonance beat placement and pitch height. When they're separated from their performative elements, do they still carry particular emotional meanings such as exhilaration or sadness? It might not be possible to answer such a question directly, but we can approach the answer indirectly by asking are there musical features that are used in conjunction with texts and are associated with particular sentiments or emotions with such regularity that we might believe that the musical features themselves are used to strengthen the effect of these affects and emotions or do they indeed carry that affect or emotion themselves.

The problem with making such connections with the traditional musical analysis is that there's no perfect correlation between any musical feature and any emotional state. We might think at least in Western common practice tradition that words connoting joy would tend to be higher in pitch than the notes around them or the notes preceding it. And it's easy to find examples to confirm this theory.

In Schubert's "Gretchen am Spinnrade" or "Gretchen at the Spinning Wheel," one of his leads every single time, the word Herz, which is the German for heart, appears or lacheln, smile, or Zauber, magic, Der Kuss, kiss, it's always at the same pitch height or higher than the preceding note and always in the upper register for the singer.

[MUSIC PLAYING]

SINGER:

[NON-ENGLISH SINGING]

PROFESSOR:

So it might sound like this is an easy connection, but then a little bit of Elvis enters into your head.

ELVIS PRESLEY:

(SINGING) But I can't help

Falling in love with you

PROFESSOR:

In this song by Elvis, originally written by George Weiss and others, called "Can't Help Falling in Love," we always hear the words love and darling below the previous note. Even in an expression as magicalistic or full of word painting as falling, even that is set with a rising motion. So this says that individual pieces, some here and some there, might have clear relationships among text and emotion and musical features. But across a large corpus, the aggregated experience of listening to common practice or folk or popular music is far more muddled.

To clarify this, we can turn to corpus studies, especially computational analysis of large repertoires of musical encoded data sets whose sheer volume makes it impossible for human researchers to analyze all these pieces manually. With computational methods for extracting connections among music, text, and emotion, we can find statistically significant trends in music affect relationships that stand out from a large corpus even if the magnitude of the trend itself is quite small.

Sophia Sun and I used a repository of lead sheets, that is melodies with lyrics and chord symbols. There were about 1,900 songs, 1,895 in total, that were in English language, that the encodings seem to be high quality and that there were no duplicates. This is a repository called Wikifonia that used to be freely available on the internet, and unfortunately their license expired and now is only available to those who had downloaded it beforehand.

In these 1,900 songs, there's about 200,000 notes and a little bit less than that in terms of the words and 70,000 chords. So this is a big repository to look at.

We also used a second corpus but this one not a musical corpus. This is an open crowdsourced database of 14,000 common English words that was made Mohammad and Turney of the National Research Council of Canada, and each of these words has been coded according to the affects that they have. For instance, words with the affect of anticipation include skewed or recreational or zeal, holiness, intimately, harbinger, birthday, something I think we all anticipate, or plan.

And we can also look at words that have that affect and that affect alone. So, for instance, if we look at disgust, we see slime, dung, barf, diarrhea, gratuitous heartworm, things that only have this particular affect. Together with a small list of what we call stopwords-- these are words that have no affect and appear too commonly such as the, and, us, she, he-- about 60% of all the words in the songs could be classified according to affect or not having affect in the stop list neutral. And there are at least 2,000 instances of each affect and some with over 12,000 that appeared in this corpus. So there were-- in the repository, there were 2,853 songs that had anger in them.

Also in the corpus were 25,000 words that were coded as being neutral. And then of those plus stopwords, there were about 12,000 words that we can be pretty sure to be neutral. And then 100-- sorry, 123,000 that could be believed to be neutral-- and then 183,000 is all of the comprehensive words that did not fill foot-- fit into one of these 10 affects.

Sun and I use music21 to automatically label every instance of the word in the data set alongside one value that we are measuring at the time. So, for instance, this is the song "Greensleeves," also known at Christmas as "What Child Is this," and in this excerpt, every word except for the word alas was identified by the composer-- by the computer as being somewhere in the data set. But only two words in this excerpt had a particular affect that had been identified. They affect-- identified love as joy and positive and wrong as negative.

So then we looked at particular values that we could measure. For instance, the one that I talked about earlier about pitch height-- so we're going to look for each word looking at the first syllable, how many semitones above or below the previous note does that word appear? So my is two semitones or a whole step above alas. Love is two semitones above my. Going forward, wrong is four semitones below me.

With measurements like this, we can test the first theory. Do positive words rise in pitch and negative words descend in pitch?

Here is our graph of the results for this particular measurement, what we call the relative pitch measurement. And what we can do on the left are the 10 affects or emotions that we were looking at-- anger, anticipation, disgust, fear, joy, negative, positive, sadness, surprise, and trust-- and we can look at words with that connotation compared to the neutral words that had been identified in the National Research Council Canada data set.

And we can see that, for instance, looking across at the top line to the highlighted-- to the highlighted figure that anger was a little bit higher-- well, there's a plus sign, so it was higher than the neutral thing. But the 0.89 is the p value. That is our confidence. How often would such a value happen by chance? And you subtract from 1 so 1 minus 0.89 whatever. It's about 10% chance it would happen by chance that all the words for anger would be higher.

And usually we don't think of p values that are below 0.05 or 5% chance or 5% chance of happening accidentally. We don't think of those as having meaning. In fact, when you do what we call multiple comparison sets where we're looking at a lot of things at once, we really want p values that are much better than 0.05, 0.01, 0.001, and so on.

And so in this chart, anger doesn't have a particular color on the NRC Mutual Set. There's no statistically significant value there, nor does anticipation, disgust or fear. But joy on the other hand is labeled with a darker blue, meaning that this is statistically significant. The three stars show how statistically significant.

So the words with joy are higher than neutral words. Negative words didn't have any sort of p value, but positive words had a slight positive correlation. They're a little bit higher, but if we go down, surprise was a strong thing. So words with surprise were higher than other words.

We can look at other parts, and, for instance, on the right side this triangular chart, we can look at how each affect correlates with another affect. So we can see that surprise is slightly statistically significantly higher than words about anger or slightly statistically significantly higher than words about negative.

The column just to the right of the yellow highlighted chart are the stop words compared to the, a, and, and every single emotional word is higher in pitch than a stop word. Just think about when I say the. It's pretty low. We don't tend to put this kind of high emphasis on stop words.

So I'm going to show some of the things that came up comparing these affects to neutral words because this seems to be the most fun and interesting part of the research. So all of this, this whole chart, will get reduced to a single column in this chart. And so we can see the relative pitch that joy, positive, and surprise are higher.

Here is the key for the chart. So I don't show how much higher or how much lower in pitch or how much more positive or negative, just the ma-- just the direction in terms of the colors. So blue colors tend to be higher in some way, and mine-- and brown colors are lower.

And then the intensity of the shade shows the p value. So the first things are shaded happen only 1% chance that they're happening by chance, and then the next shade is 1/10 of 1% Then going down 1/100 to 1%, 1,000th of 1%, 1/10,000 of 1%. And you can see on the right as 10 to the negative 5, negative 6, negative 7, you'll see some of the charts were really showing things that are absolutely positively not by chance, things with 10 to the negative 200ths or more.

One of the things that we thought was really interesting, we looked at the beat strength of every note. So we compared to the prevailing time signature, which is quite often 4/4, and we were able to figure out what's on the beat and what's off the beat. And one of the things we can see looking at the blue parts that words connoting joy and positive emotions are really much more common to happen on the beat, similarly with surprise and trust. But looking at things that happen off the beat or on weak beats, that's where you put the words for anger or syncopate, disgust, or fear or negative emotions tend to happen on weaker beats.

Similarly, we can look at the length of notes that almost nothing has a positive correlation that is to they say that is longer than you would expect. But these negative emotions-- anger, disgust, fear, and just the general negative-- tend to be on shortened notes so that we're really just spitting out these short emotions.

Intriguingly, when we look at the overall pitch, the pitch of a note compared to the average pitch of the entire song, there's basically no correlation. We don't tend to sing happy things high and sad things low. Only a little bit with positive happy, joy, surprise-ly things do we set them a little bit higher than what was happening before.

Did look at consonants. What notes with-- what words are set to notes that are part of a consonant chord? Remember these are lead sheets, so we have the melody, the words, and also the underlying chords. We can see some of these negative emotions, anger but especially sadness and some are on more dissonant thing notes-- than the dissonant chords than the prevailing piece but also somewhat surprising joy. Maybe that's the way you emphasize joyful words or surprising words that you put them on slightly more dissonant words.

Looked at two ways of classifying the mode of a piece, so major or minor, and we looked at do certain emotions appear more in songs that are in minor versus in major? And this is something that I at least had always been taught and I think many of us have been taught that minor pieces are sad. And maybe minor pieces are sad, but the words aren't any sadder than major mode pieces. There's no distinction between major and minor when we look at the whole piece.

So instead we used-- oh, and by the way, we will be looking later in this semester and how we can computationally measure the key of a piece when we get to Caroline-- Carol Krumhansl's work. So we thought maybe it's not the whole piece, but maybe that these sad emotions appear in a minor section or piece using a lot of minor chords within a larger piece. So we looked at the local mode, just the mode of the four measures on either side of-- I think it was four measures-- on either side of a word.

And there we found a very small correlation, but it was exactly the opposite of what we expected. In a local section, we use brown to represent minor in this case and blue to represent major. Joy and positive emotions tend to appear a little bit more in locally minor pieces.

What it might be is that there's not this connotation of sadness or negative or anger with minor, but about not knowing where you are. So we looked at the tonal certainty. How clear is it what key we're in? And here we found some correlations that were negative. That is to say, anger and negative words tend to appear a little bit more in tonally uncertain passages. Now this isn't necessarily a guide to composers for what to do but what composers have done, what people have done in the past.

So what did we learn by doing all of this? Words connoting joy and surprise are higher in pitch than their surrounding context, but negative connotations such as anger or disgust are not lower in pitch than in neutral were. Positive words are not highlighted with consonances. If anything, it's slightly the opposite. In one measure of dissonance, however, words connoting sadness are more common, but negative words appear more often in passages of tonal uncertainty.

Finally, mode has almost no correlation with the emotion of words whether we're talking about the mode of the piece or whether the chord at that moment is major or minor though surprise and anticipation do get minor chords more often.

It may be that this connection of minor with sad is one of these cases of a purely musical connection. That is to say we've learned to hear minor music as sad regardless of what the words are about.

And this is something that, as a music historian I know hasn't been true all throughout music's history. There were times when a pure major chord was considered harsh or a little bit sad, and that's not the case today. But maybe when people do consistently report in music psychology surveys that they hear sadness in minor pieces that there's something happening that's not fully explained in the words. So there's always more research to happen. And I just want to thank you for allowing me to share what Sophia Sun and myself have found.