

[SQUEAKING]

[RUSTLING]

[CLICKING]

MICHAEL Let's get started. So we're going to number the classes as if when you have a substantial thing to do over the weekend, as you did this time, that that counts as a class when you watch me talk on a screen.

So we'll call this class IV, even though this is the third time we get together. And last time we were together, I suppose this should be penultimate class, if I'm numbering the other one, we talked quite a bit about what is a note. And let's just take a second to review because there were a couple people who weren't here on Friday. What are some of the things that are complicated when you're talking about what is a note? Quick hand and shout when it's quiet or just--

AUDIENCE: Does it need a pitch? Does it need a duration?

MICHAEL Does it need a pitch? Does it need a duration? Great. What was the problem when we went [VOCALIZING].
CUTHBERT:

AUDIENCE: You couldn't tell if notes were slurred.

MICHAEL You couldn't tell if notes-- not so much slurred. You can tell-- what's the other word he's looking for? Tied? Yeah.
CUTHBERT: Because [VOCALIZING]. That could be for tied notes, depending on if we think of the notation as one domain-- as our definition of the note, or if we think of what we hear as a domain. Good. Anything else that's screwed up our easy conception of what a note is? Remember or make up your own. We had 14 notes in a glissando. We had-- yeah, go ahead.

AUDIENCE: I forget what it was called, but the term were you like--

MICHAEL Well, sing it for us.
CUTHBERT:

AUDIENCE: You have two notes, and then you have to oscillate between the two.

MICHAEL Well, what was that called? Somebody else? Yeah, John.
CUTHBERT:

AUDIENCE: Trill.

MICHAEL A trill. Yep, trill is one of them. And the other one that we talked about was? Yep.
CUTHBERT:

AUDIENCE: Tremolo.

MICHAEL Tremolo. How many notes does it take to write a trill on a standard score?
CUTHBERT:

AUDIENCE: One.

MICHAEL One. How many notes does it take to write a tremolo?

CUTHBERT:

AUDIENCE: Two.

MICHAEL Two. How many notes are there in a trill? Whole gob, a lot them, as many as you can put in. So these kinds of
CUTHBERT: different representations we did a lot on. Today we're going to be talking more about not just what is a note, but how can we represent a note or how can we represent something else for a computer.

You know what? I made really nice handouts that I put online and forgot to put them out. So I'm going to keep one for myself, and we'll try to do a group here going this way, a group here going this way, and in the back, where I see people the least, I see the fewest name tags. So if you don't mind, I'd really like to get to know everybody by name.

There's some discussion of various configuration problems and a few solutions and a few things still left unsolved on Canvas. If you'd like, stick around after class, and we'll talk through some of those things. And there was another question-- who raised it? Hannah? No. About whether or not the standard library is OK. So the standard library-- and that includes things like math and regex and things in Python that some classes don't let you import.

Yes, if it comes with Python, when you download it, you may use it on anything. So I'll unlock the entire standard library. So go ahead and import math within the math library. Some people saw that there's a particularly useful function in the math library. What's that?

AUDIENCE: Log.

MICHAEL Log. When you're dealing with frequencies and things, you'll see the logarithm function. So if anybody
CUTHBERT: reimplemented the entire logarithm function, we can talk about getting a life-- no, no, no, you did great. Super. We can give extra points on that. But if it's something that you have to download via pip, then always ask beforehand. But I do want to say, at the end of the semester when you have final projects, it's anything goes. So if you can find it on the internet or you can build it. You can use it.

We're going to start-- we're not as talkative as we were on Friday. Friday, we were very, very energetic. So maybe we'll start with talking amongst ourselves. I want you to-- you have one byte. How many bits are in a byte?

AUDIENCE: Eight.

MICHAEL Eight. Good. So you have eight positions where you can have ones or zeros. And I want you to try to encode pitch
CUTHBERT: and duration. So the two things that in one representation, one definition of a note, are two of the fundamental parts of it. I want you to try to encode a note in one byte and then talk about what kinds of things you've had to decide you're not going to represent in order to make that possible.

I'd like you to do this in groups of two. So it's very-- yeah, everybody is either at a desk with somebody else or not at a desk with somebody else. So if you could just make that the group, and we'll give it about three minutes. Let's come together and let's see.

We started middle of the alphabet last time. Let's choose a group involving somebody with a name toward the end of the alphabet. I don't think we have any Zs, Xs, Ys. Vanessa and Adam. Either way, I had started this. You guys would be going first. What's something that you decided to encode and what's something you decided to compromise on?

AUDIENCE: I think we decided that it was more important to preserve information about pitch than it was to preserve information of duration.

MICHAEL
CUTHBERT: So that's always a choice that-- so we're going to have a lot more information or all the information on pitch rather than on duration. Good. So we have to make a compromise that, even though I said these two things, maybe one of them is more important than another. Great. Going alphabetically backwards. Oh, we should have started with Vincent and Jake. Sorry.

AUDIENCE: I think what we decided first was that, like Vanessa's group, pitch is most important, and so we decided to prioritize encoding the 12 discrete chromatic Western pitches.

MICHAEL
CUTHBERT: So we've decided to limit to Western pitches, which is already a compromise, losing all those great microtones. And we said we want 12 of them. Great. How many bits is it going to take to encode 12 distinct things? Let's do fingers. Good. I see a lot of fours because-- why is it four? Somebody explain. Anyone want to-- yeah, Jordan.

AUDIENCE: Two to the 4 is 16?

MICHAEL
CUTHBERT: So 2 to the 4 is 16. And what's I have to do with 12?

AUDIENCE: You can have any combinations, I guess, permutations.

MICHAEL
CUTHBERT: You can have that many permutations into there. What's 2 to the third? Everybody.

AUDIENCE: 8.

MICHAEL
CUTHBERT: 8. So it's that 8 is less than or-- well, no is, less than 12, was less than or equal to 16. Therefore, we need 2 to the fourth. So we need four bits for that. So we have four bits in one thing for pitch class. Is that a term that everybody knows? It's not a term taught in all of 301 first semester theory.

Well' I'll just say encoding C as 0, C sharp or D flat as 1, D as 2, or also I suppose C double sharp or E double flat. We don't really care. So we're going to use four things for pitch class, something like that. Great. And anybody else do something like this? Any other groups? Jason and-- OK, so quite a number of groups. Let's see. We'll go with Phil and Angelica's group. What did you do with your other part?

AUDIENCE: Oh, we try to split it up with duration. So it's twice the octave pitches, just take over the duration of the rest of it.

MICHAEL
CUTHBERT: So how many durations can you encode with what's remaining?

AUDIENCE: 16?

MICHAEL CUTHBERT: 16. What is some durations that you might have chosen to encode? Or we'll move on to a Lilla, Lila and John.

AUDIENCE: We could do whole notes, half notes, dotted quarter notes, et cetera.

MICHAEL CUTHBERT: Whole notes, half notes, dotted, quarter notes. Explain a little bit more of your et cetera.

AUDIENCE: Oh, yeah. So you could have your best whole, half, quarter then either dotted because we have 16 bits to work with.

MICHAEL CUTHBERT: OK. Whole, half, quarter, and--

AUDIENCE: Eighth.

MICHAEL CUTHBERT: Eighth. And so how many bits is it going to take if we just stop there?

AUDIENCE: Only two.

MICHAEL CUTHBERT: Only two because we have 2 to the-- 2 squared is 4. So we can put that into there. And then if we want to encode if it has a dot or note, is that another bit? So then we have dot or not. What happens if we say, huh, we just put one more thing in, we'll put in 16th? What do we get as a problem? Can we squeeze that into two bits? No.

So once you decide that you're going to be able to encode 16th, you might as well go and do your 32nd, 64th. Do we want to do 1/28 or breve, double whole note? They're both pretty uncommon. So we're slightly wasting some space on some pretty uncommon things, but we get the dot. So that's three-- oh, yeah?

AUDIENCE: I was thinking about doing just 16th notes as your base, and then you can have all multiples of 16.

MICHAEL CUTHBERT: OK. You can have multiples of 16th. So you have 16th as a base and then--

AUDIENCE: If you had the number two or one, I guess, you'd get-- like 0 would be 1/16 note-- or I guess-- or maybe not. Maybe you'd want like-- maybe you'd want to be able to have zero time, but if you had 0 as 1/16 note, then 1 would be-- 2/16th notes would be an eighth note. And then you could do like a dots and stuff, and you can do all of that stuff.

MICHAEL CUTHBERT: And you can encode up to something that has 16/16th notes, with that way.

AUDIENCE: Yeah, would be a whole note.

MICHAEL CUTHBERT: Which would be a whole note. So that's pretty great way. It does mean we are spending quite a bit of energy dealing with things that are a little bit unusual, but yeah-- anybody have any ways that we didn't talk about very much in this-- yeah.

AUDIENCE: We thought that it was important to encode octave information too. So we spent six bits encoding the 64-- or 64 of the MIDI note numbers and then two bits on duration.

MICHAEL 64 MIDI notes. So now we have octave with 64 out of the 88 keys in the keyboard. You decided to do just the outside ones and leave out the ones in the middle. OK, good. And then two bits for duration. So we get maybe whole half, quarter eighth, not too much more. Great. Super. This is great thinking exercise.

CUTHBERT:

Why you want to do that? Well, there are some people who use-- oops, that's the most confusing slide in a while. There are quite a number of people who are very interested in not even one byte but 6-bit encodings. And for instance, in the Braille music code, one of the important things is being able to encode the seven diatonic notes because the diatonic notes tend to be used much more often-- and maybe somewhere else, you have a sign that tells you the key signature-- the seven diatonic notes and eight durations in six bits.

How do they do that? How do they get eight durations? They've doubled up and said, look, if you're on 128th note, you're really unlikely to move to a whole note. So that contextually, you take advantage of the fact that it's very rare to have a whole note and then a 64th note. And if it's really, really, really unclear, fortunately, they put in a space, the Braille music code, puts in a space where there's an end of a measure.

So if you're really unclear, like, I wonder if these are all 128th, no, no, no, the measure ended. They're probably the long notes. So that's one reason why you might be thinking of very efficient encoding. Many things are lost there.

I don't want to spend as much time on this but get back in your same group and talk about what you could do if now you had three bytes on average to encode your notes. So that if we have a typical score-- and I'll leave you to think what a typical score looks like-- that you could be doing-- sometimes you take a little bit more, sometimes a little bit less. So just talk through. You don't need to come up with a full representation like this for it, but just-- what are some of the things that you might add?

OK, let's come back together. What are some things we use this enormous bounty of space of three bytes to do that you didn't do before? Let's just go around, and you can say pass, but we'll start over here in second row and go, I guess, left to right.

AUDIENCE: I would say [INAUDIBLE] for the [INAUDIBLE]

MICHAEL OK, great. We've stopped leaving out octave for everybody. Another group.

CUTHBERT:

AUDIENCE: Dots.

MICHAEL Dots. We have dots in here. Great.

CUTHBERT:

AUDIENCE: Better articulation.

MICHAEL Articulation. Ooh, yeah. Does every note need an articulation? No. So this is where average might come into play.

CUTHBERT: You might have something-- some sort of a flag somewhere, some condition that says, more information to come, and then you can use the additional space. Great. Articulations. Other things?

AUDIENCE: We said maybe dynamics or volume.

MICHAEL Dynamics or volume-- --mics/volume. Just in case I forget to mention it at some point, a lot of people in the

CUTHBERT: computer music world call the dynamics of volume the velocity, which really tells you that the piano was the first instrument they were thinking of encoding, because it's how fast you hit the note-- how fast you hit the key on the clarinet doesn't really make it any louder. But velocity-- dynamics, volume, velocity. Good. Continuing on.

AUDIENCE: I guess you could do something like what you just said for the flag, where if it's an extreme range of a note or extreme duration, then you say, oh, this is an unusually long note.

MICHAEL Extreme cases to follow or something like that, to follow and-- so that could be-- these could be little flags,

CUTHBERT: suppose. Do people think flags look-- flags look better going left, don't they? Good. Other-- continuing around, other things chose that aren't on here?

AUDIENCE: So I think one thing was, I guess, time signature and stuff because it encoded with just the reference of musical half tone and quarter notes. Information isn't necessarily stored within the--

MICHAEL I actually made-- I put it in, and anything else on this slide. No, I didn't put it on the other one because I thought

CUTHBERT: maybe time signature could go in, key signature. Other things to go at the beginning. What else comes at the beginning piece? Just shout out.

AUDIENCE: Key signature.

MICHAEL Key signature. Time signature.

CUTHBERT:

AUDIENCE: Tempo.

MICHAEL Tempo.

CUTHBERT:

AUDIENCE: Name of the piece.

MICHAEL Name of the piece. Metadata. Name of the piece. Metadata. Very important. Who wrote it? Metadata. Sorry that

CUTHBERT: was yesterday's, metadata. Now metadata. Clef. Good. Quickly, go through the front row if you have something that you'd like to throw onto this. We'll continue. Adam, you guys just spoke, right? No.

AUDIENCE: What?

MICHAEL No. Because you shouted out when I said, anything else you put in your encoding? Vanessa. Vanessa.

CUTHBERT:

AUDIENCE: [INAUDIBLE] whether or not [INAUDIBLE].

MICHAEL Ties. Good. That lets you build a lot of things. Jordan, Misha, anything?

CUTHBERT:

AUDIENCE: Triplets.

MICHAEL Triplets. Yep. Triplets. I was just-- what's the generalization of triplets if they're not three? Anyone know term for--

CUTHBERT:

AUDIENCE: Tuplets?

MICHAEL Tuplets or-- tuplets or tup-- more often tuplets in the British pronunciation. Tuplets more often in the English pronunciation. Good. Anything else? Why would we want-- and so here are some examples, by the way, of how the Braille code uses an average of one byte but allows for-- or one cell, six bits, but allows for other things.

CUTHBERT: So if you have accidentals, we can be adding particular marks to say, hey, there's going to be something else coming next. Look up what this extended mark is. Or sometimes, there is a sign that you can use for saying, hey, I know you think that I'm all in 30 second notes. I'm probably going to jump to a 16th note, but actually, I'm going to jump all the way up to a half note or something. So there are some cautionary signs for that. And beams and things of that sort can be put in. Yeah, Adam.

AUDIENCE: Is there a Morse code for music notation?

MICHAEL Morse code? Not that I know of. Not that I know. The Braille, by the way, music code predates the Braille text.

CUTHBERT: Louis Braille invented the music code first and then decided to do it with text before, in part because text is already a little bit easier to read with raised impressions, whereas music was just impossible to read with raised impressions. So there was more of a need for reading music than there was for text at the time.

And just for a second, let's just think, what could you do that you couldn't do with this if you had about a megabyte per note? That's a million bytes. People have thought of this all the time, and you've probably done it, and you're not even thinking of it. You've probably used a megabyte encoding per note already today. I know that John has. Just--

AUDIENCE: So I see different features for-- like the output of the notes or whether it's like a stereo or mono sound.

MICHAEL Stereo or mono sound or-- you're getting-- yeah.

CUTHBERT:

AUDIENCE: Fidelity.

MICHAEL The fidelity. Just the audio wave of the note. You can encode the exact-- why do you need to encode-- is this a C sharp or something when you can just hear it? And you can also be encoding without the fidelity, the mono thing.

CUTHBERT: You're encoding the exact instrument, how it was performed.

You're losing generality because you're encoding one performance of a particular piece and not every performance, but maybe you can do one megabyte plus an average of three bytes to also store what the note was supposed to be. So we'll be able to get the best of all worlds.

Why do I care about efficiency? It's in part because we're at a point where we can start storing a lot of information about notes. Most of us don't bat an eye anymore at the notion that we're streaming in high def or 4K or whatever and gigabytes of billions of-- did I say billions before? I meant millions before. Billions of bytes in an hour just to watch a movie.

So why do we care about if we can store one extra thing in the note? It's because almost entirely based on history, this is one of the earliest encodings of a musical piece, the dissertation from 1982, but this is going in the '60s and things. And we started with punch cards. And punching cards takes a long time. So the more efficient you can get with your encoding, the fewer cards you're going to need. And this has left consequences on us for today.

By the way, I always thought that, oh, punch card, you get 1, 2, 3, 4, 5, 6, 7, 8, 9, 10. You get 10 bits of information. It turns out you punch too many holes in a punch card, and you try to put it through the punch card reader, and it jams. So you actually have a limited number of ones you can do per byte or per row in a punch card. So we're still having to deal with some of the choices that were made in order-- back when computers might have 4 kilobytes of memory or 16 kilobytes of memory for every single thing.

So if you want to store 16 kilobytes or 16,000 bytes, if you're storing 10 bytes per note, you only store 1,600 notes. You can't get through a whole piece. But if you can get it down to two, matters. And so we have a lot of ways that are very compact representations that are still out there and that we're only over time feeling like, you know, we can expand a little bit.

This is your first encoding format that you've seen in this class, I believe, besides stupid, weird ones I make up on the fly. This is called MuseData. This was one of the first major projects of encoding a lot of music into things made by Walter Hewlett of Hewlett-Packard fame. And let's try to guess what some of the things might be in this. Let's start with something really obvious. What do you think this means? Rest. Good.

We're obviously already in a period where we can afford 4 bytes to write out rest rather than putting R in. This tells you that this is a little bit later than other things that were happening. Somebody else take a guess what something else might be. Yeah, go ahead, Joe.

AUDIENCE: Does q stand for quarter?

MICHAEL CUTHBERT: Q might stand for quarter. Let's see. We might find something that will help us with that. I'll go easy. Measure stands for measure, and measure stands for measure, and 1 stands for 1, and 2 stands for 2. Is there anything that would let you think that q, q cube in one measure might be the right thing? Yeah.

AUDIENCE: T, 3 times signature probably?

MICHAEL CUTHBERT: T time signature 3/4. Everybody's notation teacher always told you never put a slash in a time signature when you're writing on the page. Well, in computer stuff, we always do because we have nothing better to separate them. Good. Q's that. I think these are?

AUDIENCE: Notes.

MICHAEL CUTHBERT: Notes. Good, good. Interesting thing is that they end up encoding the duration twice since this one, I've tried to do a long time ago. What do you think that is? And it's almost impossible to get. So you have the quarter note, and you also have something that says a quarter takes 2 something, takes. Maybe some-- things like that. And so you encode it twice. Once-- what's the difference between these two-- each quarter is two of something, and that this is a q. Who needs what for what? Yeah.

AUDIENCE: What's printed on the page versus what's played.

MICHAEL CUTHBERT: What's printed on the page and what's played. So that you don't need to convert between visual and time duration. And this is a function in almost all modern storage things still today. It's also helpful for-- there are some very weird cases that you'll run into in one of the readings later, where composers have used durations to mean something slightly different from what they generally are.

So this is some of the reasons why we went through efficiency. You're going to be-- on the second problem set, the one that you'll get on Friday, which will be done in groups, you're going to be coming up with your own representation system for some sort of thing that is music, except for the thing I'm just about to talk about, common Western music notation, so something else besides normal scores that we see all the time. And you'll need to make the same choices, decisions, mistakes, maybe that pioneering people did in the '60s and '70s and hopefully learn something about what trade offs happen along the way.

So we're going to be talking-- oops, slides a little bit out of order. This should be here. This slide should be-- about something in the class, which is called common Western music notation, often abbreviated to leave out one letter, either common Western music notation, common music notation, common Western notation, some things like that. This is an example of what, I think, most people would believe is common Western music notation. It uses notes, staves, clefs, things. Left to right means earlier to later. Up to down means

AUDIENCE: Higher to lower.

MICHAEL CUTHBERT: Higher to lower. So that here's a whole bunch of high notes, and here's a whole bunch of low notes. Now, except for that there's certain things that we say only within this particular staff is higher and lower, and then we have divisions, and we move on to other systems and stuff.

So these are things that are standard in common Western music notation. That was a piece by Joseph Bologne, the Chevalier de Saint-Georges. This is a piece by Clara Schumann, her piano trio, trio for piano, violin, and cello. And you see some things that we consider to be common Western music notation, including slurs, ties, grace notes, beams, all these types of things.

Common Western music notation can represent a piece of music in multiple ways. This is an aria from Wolfgang Mozart's opera Don Giovanni, an aria of song, in that [NON-ENGLISH SPEECH] by Donna Elvira. And we have the top of each of these systems, voice part, Italian, translation to English, to two staves for piano, and then the exact same music in some way. Here's her part.

Here, except for the English is gone, very helpfully replaced with German. And her nice treble clef was nicely replaced by soprano clef for ease of reading for all of us because she's a soprano. And the piano is gone, replaced with-- take a guess. What are these parts for?

AUDIENCE: Orchestra.

MICHAEL CUTHBERT: Orchestra. We have strings. We have other things going on. One of the questions would be, when we have multiple representations of the same piece, how do we translate from one to another? Is there a way that we can go from this exact moment to that exact moment, and what kinds of relationships would make that possible?

So jumping back to Clara Schumann's piano trio, there are some things that violate particular norms of common Western music notation, such as-- well, they're part of it, but they violate the basic, basic principles. What are these things up here?

AUDIENCE: Grace notes?

AUDIENCE: Grace notes.

MICHAEL
CUTHBERT: Grace notes. Let's say a quarter note takes one tick. A half note takes two ticks of some sort. An 8th note takes half a tick. How much does a grace note take? Yeah.

AUDIENCE: Usually, half the allotted time.

MICHAEL
CUTHBERT: Usually, half the allotted time? So if I put-- that's very true in one sense. So correct. Absolutely correct. It's also not true in another sense. So if I have 4/4, and I put in a dotted half note, and then I put in a grace note, that takes half of that. So that's 1/4. So I still have enough room for that amount of space. Can everyone see from back here that's a dotted half note. What's wrong with this representation? Adam.

AUDIENCE: The grace note also eats away at the actual time for the note.

MICHAEL
CUTHBERT: Yes. So the grace note steals time from the note that it's next to, or it steals it sometimes from the one before. So in terms of-- how many grace notes can you-- how many Grace notes can you put in a 4 4 measure? Let's put it this way. It's a trick question?

AUDIENCE: 16.

MICHAEL
CUTHBERT: 16, no. You can put as many in as you want because the grace notes just steal time. So for instance, here, we're in 6/8. So we have two dotted 8th note bits, and we have-- a dotted quarter note. So we have a dotted quarter note here and the equivalent of a dotted quarter note here. And we can squeeze in five grace notes there, seven, eight. Doesn't matter.

But where conceptually do you play that? And if we have this notion, that in common Western music, everything that happens at the same time is vertically aligned-- I hope your theory teachers taught you that. That's the most important rule in general-- how do we know where to align these notes with the grace notes? So that's one violation thing.

The limits of common Western music notation are not strictly defined. This is a very common genre of music from around 1700, 1687. Élisabeth-Claude Jacquet De La Guerre, Jacquet De La Guerre. Very interesting piece. A lot of things that look like common music notation until you start getting-- what is this? What are these?

All the symbols of common Western music notation but used in a different way. These are unmeasured preludes, so certain places that should be played in time and certain other things that most people believe these tell you play these notes for exactly how long you want. I kind of figured some other people think maybe improvise something on those notes, but--

It's common Western music notation, but it becomes a problem for the computer. What's the problem for the computer? When it's all handwritten. And so there's all kinds of ambiguity. You can tell that, OK, this is going, going. and that's an A and that's-- that's an and that's an E and that's a C and-- oh, that's either a D or an E. We can conceive of a note that is hard to read. Does your representation have a way of notating that?

So handwritten scores also tend to contain lots of things that violate common music notation. I saw a Beethoven manuscript once where he was just going along, going along on two staves, then it became obvious that the top part was just going to repeat-- sorry-- the bottom part was just going to repeat what he was doing. So it goes along and just uses the top one. Next time, he's drawing the next system, and he needs another staff. Oh, and there's this empty staff above. So he goes ahead and uses it, so you can be jumping forward and backwards in time in the same thing.

There are small-- almost every composer of significance has made their own tweaks to common music notation. This is by the 20th century composer Igor Stravinsky. Later in his life, he started to ask, why do we-- when you have instruments that don't play very much, why do we waste all the staff putting them out? So it's really easy to see what's resting if you just start the staff where they come in and end it when they end. So a little bit there.

Limits of common Western music notation come out in certain pieces. This was, I think, my freshman composition assignment when I was trying to be a composer, where I've used some things that look like normal notes, but then at this point, I didn't want to calculate what the tempo was.

So just hold it for about two seconds, hold this for about two seconds, and then get to this point. And this part should accelerando, so we'll just show about how end like a 64th note. Is that part of common Western music notation? It wasn't 50 years ago, but there's quite a number of scores that do it now.

A lot of the music being written today are stretching common Western music notation quite to the limits. This is Kate Soper, who was a Pulitzer Prize finalist a few years back, teaches at Smith College. One of her pieces from part of *Ipsa Dixit*, her opera. She herself set it, where moves back and forth in the voice between sung and spoken and uses these other symbols to tell you how much to open and close the mouth.

And something a little bit odd about the violin. Lots of things odd. I'm thinking about the staff itself. It's not a usual clef for the violin. No, violin in percussion clef with four lines. What do you think those four lines represent? What does the violin have four of?

AUDIENCE: Strings.

MICHAEL CUTHBERT: Strings. So just what string to play on. You'll figure out the notes. This is a piece by Brian Ferneyhough, part of a compositional group called New Complexity, which, if we look just at the first measure, there's almost nothing that violates common Western music notation principles that we learn in our theory classes, but they tend to take it to the extreme. So for instance, the first measure is 2/8, where we're going to put in place of four-- what? The 4/16th notes of 2/8, we'll put 7/16th notes in.

Within that, for most of these, we're going to put 21 30 second notes in the place of the first 20 of them, or 24. And then within that, we can have triplets so that tuplets can be nested almost indefinitely. And there's no reason why our time signatures have to be powers of 2. So people are playing with things a lot.

OK, Professor, stop talking about the weird music. This is all strange new music. Nobody in the great past like Mozart would do things like this. We go back to Mozart. This is, again, *Don Giovanni*, the end of the first act. Very cool moment if you don't know. The opera, they go to a very, very rich ball, and it's kind of like going to a club today with multiple rooms and multiple people, DJs, or whatever, doing things.

And so here at this big ball, we're going to have in one room in 3/4, a stately minuet for the upper classes, like the 3/4 dances at the time. And then in another room, we're going to have a dance in 2/4 for the middle class because they like those things. And then for the farmers and things, we're going to have-- they like their triple time 3/8, 6/8 dances all together at the same time, the bar lines.

When you say, well, let's start at measure 22, you go I'm sorry, my measure 22 or her measure 22? Because you don't have a concept that measures-- are universal. I think this is-- yeah, this is Ravel or Debussy. I think it's Debussy, *La danse de Puck*. Very, very common. Everything looks great. Everything looks great until-- what does that mean? We're in treble clef and treble clef, but we have a bass clef just for this one note. Does it also go here? No.

You can tell what this is, what that means. Just better not to write a whole other staff just for one note, but you can have two clefs simultaneously on the same staff. There are examples from Bach of two time signatures in the same staff or changing time signatures mid measure. And what does that mean? So almost every composer has pushed the limits in some ways of conventional music notation.

Jump back to Clara Schumann one more time. There are certain things that might be very obvious to us, but they can be very difficult for encoding in such a way that it's completely unambiguous for a device as stupid as a computer, such as, where in the measure does this begin?

You all are musicians. You all know that this line is continuing up to here, and that's why there's something empty here and empty here. But how would you teach a computer to make that kind of abstraction, that lines can move from staff to staff and therefore don't need to be filled in? So any questions on this-- sorry-- on any of these things or any discussion points. And Adam, we'll start with you.

AUDIENCE: What is that rest supposed to mean right there?

MICHAEL What is this rest?

CUTHBERT:

AUDIENCE: Is at the beginning of the [INAUDIBLE].

MICHAEL Oh, this one here. This is because your professor did not print a PDF but took a really fast screenshot from finale.

CUTHBERT: And so it's even a new thing, that-- our user interface. So this is the user interface that says, this is a hidden rest until we get to this note that has to be added onto here. Super great. No one's ever seen that. But yes, the user interfaces also end up being part of our conception of notation today. Good. Other points or--

So where does common Western music notation end and where does other things go? And I'm sorry-- and what might we want to represent that's not common Western music notation? So this is a type of thing I work on from time to time. This is church music from, I think, the early 17th century or late-- sometime, around late 17th century maybe. And looks very much like ours, but just the note shapes are a little bit different.

Common Western music notation, just with a little twist, maybe. Not sure, but it's not too far off. All of our notation goes back and back and back in time, or it has come from things, decisions we made before. Let's start with-- this is Western music notation. This is the first music notation of Western Europe from about the ninth century. This is the chant for the mass on Christmas Day.

And what these things do is they tell you how many notes there are and if they go, in this case, 1, 2, 3. And they go-- start low, go high, come back down, about the same, about the same. Higher note, same thing. We have up, down. Most of these go up and down. Here we have 1, 2, 3 notes going down, three notes going down. The height means nothing. It is music. Might we want to represent this and how might we want to?

And some of the things people are always wondering, we have pages and pages and pages, thousands of pages of music in this notation. How was it helpful for somebody to know how many notes to sing and whether they go up and down but not what notes they are? So that's something a lot of people got. Wish we had met yesterday. Would be a little bit more obvious.

This is a piece from around 1420 about Cordier's piece, *Belle, Bonne, Sage*, in the shape of a heart. There's something about the shape of the piece, is that something we want to represent when we're making our representations. About the same period, it's a piece in the shape, instead of a heart, in the shape of a harp-- hard to say-- where it's basically the music notation of its time, except that it's sideways. But we're writing it on the strings, so we need twice as many strings, and we don't write anything on the spaces.

Getting back to more modern music, some of the things that definitely violate almost every principle of common Western music notation but use some of them. This is a Cornelius Cardew piece treatise, which uses a lot of notational symbols, a lot of things that look like music notation, but it's up to you to interpret them. Done differently. Every group performs this quite differently.

If you're going to have a graphical score where you use music notation elements but they don't mean anything, maybe a musical score doesn't necessarily need to have any elements at all. Earl Brown's folio from December-- yeah, December, what is it? 1962. That's all things that have come out of the West, Japan, China, every--

Many cultures around the world have developed their own notational systems, and then there are notational systems that have been made to try to study the physicality, in this case, of Inuit throat songs of-- without using anything that might have the cultural baggage of Western traditions. So we start notating the spectrograms. Is that a form of music notation? And if so, how would we do that, one more of tablature? Just giving a sense of some of the notations that might be happening in the world.

In addition to notations that have been sort of created over time by people, there are invented notations that tend to-- that have been found to be very useful times. This is something called TUBS, if anybody has taken-- the Time Unit Box System.

I think it was invented around the 1970s. I'll look up the exact time later, but the notion is kind of like I did with the bits, things that you can specify what the smallest tick you want to notate is and specify when somebody plays. Or if you need to specify different things they might play, you can put that in the box. This has become a very useful and used system, especially in the notation of West African music.

Lots of notational system representations have been created, many of which have not been successful. This is the turntablist notation system, because there's nothing that's greater and cooler than spinning at a bar while reading music notation. But it is a way of notating what somebody did for reproduction later. Let me see if my sound is working before doing this. I'll be very disappointing. I should have checked that ahead of time. Give me one second.

[ROCK MUSIC] Grab it. Grab it. Grab it, grab it, grab it. Grab it, grab, grab it. What is it?

But not all notation and representational systems end up being very successful. I think this is super cool. I don't know why I made a joke about it. I think it's really innovative. The website has been down for about six years now, so I assume it did not take off.

One of the earliest mechanical notational representation systems that's still very influential is the player piano roll, or we just call it the piano roll now, where the actual punching of the notes on the staves becomes a very important part of how-- well, it encodes the notes.

So we get, again, the same kinds of notions of up and down, the side and side, for height, I guess, height turned sideways. And as the roll progresses, there's a cursor-- there's a technical name for it. I can't remember it in piano rolls-- that is reading at that moment what should be happening or what is happening at that moment.

There are other ways of controlling-- usually on the side what the dynamic of the entire piano is, but usually not of encoding what the dynamic of any individual note is. And quite often, written along the side, there will be the lyrics of the song written, kind of like when you're driving and it says like, slow, pedestrian crossing, read upside down, so you learn to read it that way. So that you can sing along with what the piano is playing. Very cool multimedia.

Questions about things going forward or discussion or just things that you're thinking about, because you're going to be thinking about something to represent that isn't common Western music notation over the next week. Yeah, Adam.

AUDIENCE: The abstract painting that you showed at the notes, it reminds me a bit of line rider.

MICHAEL CUTHBERT: Line rider. Oh, yes, the game where you're riding on the-- you draw all the things and then you get to see how the motorcycle jumps around, is that one?

AUDIENCE: Some people have done composition in which line rider motions are synced to a piece of classical music.

MICHAEL CUTHBERT: Very very, very cool. Nice. Yeah. I want to see that at some point. Yes.

AUDIENCE: Is this just kind of like a big music box?

MICHAEL CUTHBERT: It is, but often, something that you could also attach to your home piano so that you have your home piano early on, and then later, they become-- or I might have my chronology reversed, but then they become things that become standalone objects.

But there were things that allowed people to feel more like they were playing the piano by having you pedal it and control the volume so that you could do-- and control certain levels of softness, kind of like using a Betty Crocker bake mix to bake your cake. You get a feeling that you really did something to make the music but not the really hard parts. Yeah, cool. These are interesting.

We have some of these over at the Lewis Music Library, which I'm going to be talking more about later as we go into representation things. But the Lewis Music Library, just over here, has put as one of its long-term visions to try to get all of the different technologies that have been used to encode and playback music over time.

So there's wax cylinders and piano rolls and stuff and reel to reels and trying to make it so that there's some that are very valuable and need to be preserved and some that we can just play with. And wouldn't want to have a problem, but if you did end up breaking the real tape, it's OK. We got more. So definitely, if you haven't been to Lewis Music Library, check out what they have there.

They also have a theremin if you want to play with that. And you might be using them at the end of the semester when we make some videos because they have a nice podcast video recording space for editing there.

So we've been thinking a lot about the common Western music notation in general terms. Now I want to think about a little bit computer representations. I don't why I didn't write the word computer representations. And again, these are slides that I'm borrowing from a composition package named *Belle, Bonne, Sage*. Exactly, they're named after that piece in the shape of the heart things.

So we talked about this extremely fast on day one, but I want to slow down a little bit more. The two main ways that people have chosen generally to encode common Western music notation for the computer are in terms of the container based thing, the box based. So you have a big box that might be the whole score, some kind of a thing, and then you have smaller boxes that might be parts. Within those you have measures. And quite often, voices become things inside that and then going all the way down to notes. So this representation.

The other second one that is very common is the time sliced-based container. So everything that happens at a certain moment in musical time is in its own slot. Some people would not put these in three different slots because-- or these four things are in four different slots because all of these things happen musically at time 0.

The clef-- there is no note or anything that-- there's no conceptual notion that the clef comes first and then the time signature-- then the key signature comes, then the time signature comes, then the first note, except that we have to write them in certain order on notation, but all of them are present at the very beginning of the score.

So I want to start talking about some of the reasons why you might want to use one of these and why they're all bad choices. Every choice you make is going to be a bad choice at some point. So thinking about the part-based representation, you can immediately see certain problems. Anyone know an ossia staff, what you use that for? So let me explain that. So Adam, I think we've been really great, but I want to hear a few more voices right now. Anyone ever heard that term, ossia staff? Let's see, then say, no.

It's a little staff-- pianists, violinists see it quite often, especially when you're starting-- when you're playing a concerto sometimes in a really hard part. No. Yeah, Hannah.

AUDIENCE: You said something about an alternative way to play that measure?

MICHAEL CUTHBERT: Yes, it's an alternative way to play that measure. This is a hypothetical one because I don't think that that would ever be an alternative way to play this. But quite often, if it goes too high for some people, especially for singers-- actually, singers see ossias a lot, because if something's not in your voice range, some things, they'll write different part.

So you have a staff that just appears for a minute. We said that our top level thing might be the parts. Beyond that, sometimes if we're a piano thing, you might have the part and then staff one and staff two, but here we have a staff that just comes up. And if everything is in a box and one of the lowest boxes is a measure, you have these problems.

What do you do when you want to say-- if everything's in a box, that a crescendo goes from the beginning to the next, how might you solve some of that computationally? When everything has to be in a box. And think of a box as like an array or a tuple. Yeah, go ahead.

AUDIENCE: Have like, a flag start something in one box, and then if I close it in another.

MICHAEL CUTHBERT: Yeah. So everything could have a start and an end, or certain things might. Maybe a note doesn't need to have a start and end, but you might have a flag for starts and ends. So there are some ways around it. Other possibilities? Maybe? Yeah, go ahead.

AUDIENCE: Maybe just put them a higher level up. Like they don't belong inside measures. Maybe they belong inside parts.

MICHAEL CUTHBERT: Yeah. So you move up one level. That's great. When you do that, that works quite often, but then, you have a second problem that comes up and that-- don't worry. This problem was going to come up anyhow, especially once Mozart brings his three bands to the page. But you have-- maybe here's your part, and you have your measures and things, and now you want to elevate the crescendo that went across two measures. You want to elevate it up to the next part, the next thing.

Let's assume that these are some of arrays, so we're putting commas then. Then you might have the problem. Maybe you put it here, maybe you put it here, but suddenly, the array is no longer sequential. In other words, that certain things can appear at points that are out of line. And so we [INAUDIBLE] call this the forward backwards cursor. So that we're at a certain point, but we need to look back or to look forward in order to see what's happening at that point.

But it's also a very good solution, and it's something that happens quite often. So two great solutions. Any other-- it doesn't even have to be that great. Those are the two industry most standard ways of getting around this. I couldn't plant a better thing. OK, good. That's fine.

I should have said those are two of the three great industry standard ways, and one of you would have come up with a third way and invented something new. Never say that there's no more possibilities to MIT student.

Here is one of the problems that we were talking about this earlier with a time-based system, the Grace notes, which take up no number of ticks in the time signature but of course, take up time to play and take up space on the page. How are they engraved? And they end up being one of the great problems of the time method.

So you could say, hey, I'm great. I'm not a bagpiper, so I don't use grace notes. Bagpipers use grace notes all over the time. So I can do the time-based system. Or somebody else can say, look, my piece is forte from the beginning to end. I'm never going to have a crescendo. I don't need to worry about-- I can use the part system.

One of the things this class is not just-- is not computational representation, but it's computational music theory and analysis. Eventually, we're going to be diving into our representations and doing musical analysis on them. And there's certain things that we often like to do. We said first class, who has written a parallel fifth? Almost every hand went up. That's what we call voice leading, where we're looking at how certain chords or certain voices at one moment move to something at another.

Anybody have a professor who frowned on this? Do you know-- yeah-- overlap, they sometimes call this, where the bottom note has jumped above the place where the top note had recently vacated. In certain styles of music, it's not very idiomatic.

Well, let's analyze, how would we find overlap in the part based-- in the box system, in the container system?

Well, we'd probably have-- there's one container for this part or measure, and then there's going to be one container for the up stemmed notes. That's a voice. One container for the down stem notes.

We need to know if at this moment, at this moment, this note has crossed over the previous notes in the other container. So we have to be able to jump back and forth among containers. And that, it's not very good to do. So the time slice. We want to look at everything that's happening here and then everything that's happening here. And then it becomes a problem of yes, but how does this note in the first slice relate to the second?

So for voice-leading problems, jump to the conclusion there, neither system works well. They both have problems. They both have big, conceptual difficulties that make voice leading one of the hard problems that always needs to be solved, specifically for music before any deeper analysis can happen. And it means that generally speaking, you're going to be writing a system for moving from one representation to another, from a time wise to a box wise, and how to move back and forth and have your system be able to have multiple representations. It's one of the major problems.

Last little topic for today-- and I think we're done with slides, so I can put this up-- is just some of the terms that you're going to come up-- actually, two little things. I'll always make sure because I don't to keep anybody late, I want to cover the last item on there. You will get problem set two on Friday after you turn in problem set one. It will also be a one week assignment.

There is one reading to add on. It's about, I think, 15 pages or so but-- 20 pages, but most of it's pictures-- that you'll find on Canvas. It is an article by Nicholas Cook, who's one of the main people in this field, called comparative and computational musicology. And it's one of the foundational texts in this field. So I want you to read it. It's 18, 19 years old now, so it's a little bit out of date.

You'll see some of the encoding, some of the ways of writing computer code are a little archaic now, but what it says is still really, really good. So you'll find that online, and we'll discuss that. There are certain questions to start thinking about. You don't need to write out answers to them, but maybe sketch down an answer to one or two of them so that you can be quite participatory on Friday.

And if you haven't been reading very many scholarly articles or this is your first time reading an article written for academics, I have a little six-minute optional video on how I read these articles. That maybe even if you are a master at-- great at doing this, it might still be useful to see how a professor doesn't read from the beginning to the end in one set. So if you want that, that might be particularly useful.

So we have about four minutes for the last topic, and it's the one thing that I want-- that you'll probably want to look at the outline for. I usually put the outline for you to get later, but something-- and it's about the different music representation, some of the different terms that we're going to be using a lot.

Anybody taking more than one or a couple philosophy classes and come across this term? Anyone? Philosophy class? What about in a CS class, computer science? Anybody ever see-- I'm not going to ask for a definition right now, unless you want to do it. Yeah, great.

So ontologies are sort of this-- ontology is the study, some people say, of essence, of meaning, of what it means to mean something, general. Philosophy and computer science have two slightly different definitions of ontology. So you're going to start seeing the word ontology in a lot of representations. It has to do with representation.

For philosophy-- and this is-- I mean, really philosophy. You're getting a lot about, what are the limits of something that makes something no longer something? So we did a lot of philosophical ontological discussions today and last class. What's something that we were trying to find the limits of? when something is no longer something.

AUDIENCE: A note?

MICHAEL A note. Yeah. Some of the things, when is a note one note? When is a note two notes? When is a note six notes?

CUTHBERT: We went with that. When is common Western music notation no longer can-- does that apply? I kind of did the frog boiling in hot water. Check slowly, slowly, slowly moving up until finally it's like, wait a second. Professor, that's a whole bunch of spots on a page. But we got there very, very slowly. How do we put a dividing line?

So there's a philosophical definition. And we're going to be using that one in this class because this is also a humanities class, where we're thinking about those types of things. What is the limit of something? What is the note? If a piece of music is every note and every duration that's in there, how many notes can you get wrong and still say that you're playing Beethoven's Hammerklavier Sonata? 10%, 5%. One note wrong, and it is not the piece.

These are some of the things that the philosophical definition does. We're not going to get to the second half here, but-- so just we'll end with the computer science usage, which is also very, very cool and related. And that would be, what are all of the objects-- yeah, objects, names, relationships used in a representation?

So ontology of your computer's file system might say that there are documents, and there are folders. Can documents go in folders? Yes. Can folders go in documents? No. Can documents go in documents? No. Can folders go in folders?

AUDIENCE: Yes.

MICHAEL So we're starting to make some kind of representation of something. Can a document appear twice somewhere
CUTHBERT: in your hard drive? Different operating systems? Yes, no. There are ways of doing it. So these are all the decisions that make up the systems and objects and representations in use.

We're going to start up again on Friday by talking about some of the ontologies that we've already been creating, some of the relationships, and some of the different ways-- starting to get back into programming. We haven't really done that for a little bit in class-- some of the ways that we might represent the relationships between notes and pitches and durations and so much else.