

[SQUEAKING]

[RUSTLING]

[CLICKING]

MICHAEL SCOTT ASATO CUTHBERT: OK. Welcome, everybody. Wonderful, warm February Friday here. And we're continuing with our second topic for the class, the first one being about notes, pitches, durations. And now we're getting how to represent bigger musical scores.

I'm going to start, though, by jumping back to topic one for a second. And I want to talk through a couple of things. First, great job on Pset1. And I can say it's nice when you can say that to the whole class without putting, for the most part, Pset1 was really great. No, every single Pset1 was very strong.

And that's my quick two second overview of having the computer automatically go through and just spit out every error that it detected. And I'll go through it by hand and figure out places where there might be partial credit where a mistake early on was causing cascading problems, and things of that sort. But some really, really great things happening.

You guys seem to be really strong on pitch names, on representation, on manipulating, transposing. Some really, really good work with inversions, for the most part, which is some of the most challenging parts of the problem set. And yeah, so I'd love to dwell for the next hour on everything you all did right. Give yourselves thumbs up for that.

But let's talk a little bit on a couple of the things that were challenging on the problem set, some of them that maybe I could have done better, some of them maybe we all could have just thought that one little bit harder. Let's start with the concept of pitch class. How many pitch classes are there? Let's shout it out on three. 1, 2, 3.

AUDIENCE: 12.

MICHAEL SCOTT ASATO CUTHBERT: 12. And they're, what? One to 12?

AUDIENCE: 0 to 11.

MICHAEL SCOTT ASATO CUTHBERT: 0 to 11. Great. Super. So we almost all got that. And what's a note that's pitch class zero?

AUDIENCE: C.

MICHAEL SCOTT ASATO CUTHBERT: C, good. So let's see if we know our major scales pretty well. Pitch class four?

AUDIENCE: E.

MICHAEL SCOTT ASATO CUTHBERT: E, good. So four is E. E flat?

AUDIENCE: Three.

MICHAEL SCOTT ASATO CUTHBERT: Three. E double flat?

AUDIENCE: Two.

MICHAEL SCOTT ASATO CUTHBERT: Two. Great. Triple flat?

AUDIENCE: One.

MICHAEL SCOTT ASATO CUTHBERT: Quadruple flat?

AUDIENCE: Zero.

MICHAEL SCOTT ASATO CUTHBERT: Quintuple flat?

AUDIENCE: 11.

MICHAEL SCOTT ASATO CUTHBERT: 11. So this ends up being-- this was one of the challenge, that there's a modulo operator. Trying to make sure everything is between 0 and 11 at first, when you're looking at the names. See? And so almost everybody somewhere encoded a array that looks something like this, I think. 5, 7-- am I getting this right? 9, 11. Who did something like that at some point? Yeah.

And then so if you didn't do that and you encoded all of them, then the hard part was enhar--

AUDIENCE: --monics.

MICHAEL SCOTT ASATO CUTHBERT: --monics, yeah. If you said that one is C sharp-- ooh, it's also D flat. OK, I think I've got it. But it's also very, very odd notes like B double sharp. Oops. I should do a proper double sharp and not a programmer's double sharp, right? Yeah. Things like that. That's correct, right? Yeah.

So if you did something like this, there's that weird thing that a lot of people forgot that when you get your modulo to get everything into 0, 2, 4, 5, 7, 9, 11 by stripping off the accidentals, then you start adding-- what is it? You subtract one for every flat, add one for every sharp so you can do-- you can have your computer be very smart, at first, with things that-- that's a E double sharp. My handwriting will get better. And you can do that. But remembering at the end that, oh, if B is 11, that B sharp is not 11 plus one, but 11 plus one-- what do I want to put after that?

AUDIENCE: Mod 12.

MICHAEL SCOTT ASATO CUTHBERT: Mod 12. Yep. Good. So that was one thing. The other thing that came up quite a bit-- and you're going to think that this is a pretty dirty trick for me. Let's start with this. This is pitch set one. This is pitch set two. Is there a way we-- is there an axis of inversion somewhere that will let us transform pitch? set one to pitch set two?

No. I'm seeing some no's. Matthew, you had the first head rock. How did you compute that very fast? Did you try all 132 possible axes of rotation and saw that there were none?

AUDIENCE: The way that I did it was I made sure that the intervals of one were the negatives of the intervals of the other.

MICHAEL SCOTT ASATO CUTHBERT: Great. So when we're talking about the intervals of one, the intervals pitch set one, what's the interval that we're looking at here? What do we call-- two? And what's the interval here?

AUDIENCE: Two.

MICHAEL SCOTT ASATO CUTHBERT: One. If we're thinking of chromatic, what if we're thinking of it in more of our traditional tonal space? What do we want to call this one?

AUDIENCE: Whole step.

MICHAEL SCOTT ASATO CUTHBERT: What's that?

AUDIENCE: Whole step?

MICHAEL SCOTT ASATO CUTHBERT: Whole step. Good. Another word for whole step? Usually major second, good. And here, half step or--

AUDIENCE: Minor second.

MICHAEL SCOTT ASATO CUTHBERT: --minor second. Good. So you knew that you weren't going to be able to find that. You didn't need to do all the computation. Let's do one that's a little bit harder. Is there something like that? And can we find some of axis of inversion that allows us to transform it into that? Why? It's pretty obvious. So go ahead and always take the easy participation points when you can. Trudy?

AUDIENCE: No, because one has four notes and the other has five.

MICHAEL SCOTT ASATO CUTHBERT: One has four notes and the other one has five notes. Good. I believe 15 out of the 17 problem sets turned in found that to be acceptable because you weren't comparing first the obvious thing, that they had to have the same number of notes. Which does happen quite often, that you're going along and you're thinking of here's a triad, here's a triad, here's a seventh chord, here's a triad. Is there a way that we can manipulate the first triad to be the second triad? Yes. Is there a way that we can manipulate the second triad to be the seventh chord? The four note chord? Not with the tools that we've learned so far.

So those are some things. And a very, very, very small-- this one, I think, is important because this will come up in real music everywhere. Very, very, very small problem. And this one-- this one you might hate your professor for doing it. Can Pset1 be converted to Pset2 by inversion?

Yes, because there's nothing here. You don't need to do any operation, or any operation you do to nothing can transform it into this, right? If you think about the reciprocal or the negation of an empty set, it'll be the empty set. And that one, I think I put 1/10 of 1% on because it's kind of a nasty thing.

But it was one of the cases where unlike the four note thing where a lot of times, your algorithms were returning the incorrect answer, this one was causing the programs to crash quite a bit. A couple of years ago, my automatic assessment system didn't-- I stole it from Eran Egozy. He's very generous with that. Did not handle crashes very well. Now it does, so we'll see that.

So that's a big parentheses around the fact that the P sets were great, were very, very strong. I think you should all be proud of that. But keep thinking of the corner cases. I especially liked how many test cases that I didn't write that I saw.

I mean, I was going really, really fast through the code because you all just turned it in this morning or late last night. But for the most part, I was very impressed with how much people were trying to get at the corner cases and things, and that's a good thing to feel free to use that kind of collaborative discussion. And we're going to encourage some collaborative discussion in the next problem set coming right up.

There was one other thing besides some reading and video watching that I asked you all to do, and I got some really great answers. And that was talking about various kinds of pitch representations, what you might do, and what some of the stakeholders might be. Who might want a particular pitch representation? And that might be different from somebody else.

So anybody-- I read through them. I'm going to say some things too, especially things that I found really great. But anybody find anything that you thought was interesting? Two different groups or audiences or stakeholders that might prefer very different representations of pitches? Yeah, Vanessa?

AUDIENCE: I think it's interesting to think people that are more into editing and really analyzing sound files might be a lot more interested in seeing frequency values, as opposed to if you have a musician who's reading sheet music, they look at a frequency value, doesn't mean anything to them.

MICHAEL SCOTT ASATO CUTHBERT: Yeah.

AUDIENCE: So they would want one of the letter and then octave values in order to be able to understand what you're talking about.

MICHAEL SCOTT ASATO CUTHBERT: Yeah. Yeah, totally. Nobody's going to see, I don't know, 519.8-- I'm just making these up-- 947.82 and be like, that's my most beautiful chord. I just love that. It's very hard to do on the guitar, but when you get it just right-- yeah. Yeah, so definitely different stakeholders for different places, different agendas.

Good. Others? So we have sound editors and musicians, and we can refine that either way. Thanks, Vanessa, that's great. Anybody else? I think-- yeah, Matthew?

AUDIENCE: Some people care about enharmonic spelling. And for some people, that's just excess information that's useless.

MICHAEL SCOTT ASATO CUTHBERT: Yes. The big fight on enharmonic spelling. Yeah. I just got a bug report in Music21 on Wednesday that Music21 was not properly calling this-- let's see if I can even do it now. Maybe it was here. Was not properly calling this a dominant seventh chord, and was saying it was something else.

But if we flip this-- here. Actually, maybe-- I hope I actually wrote a dominant seventh sign. That sounds like a dominant seventh chord. That's right, right? I'm not too far away from teaching music theory for a year or so. Yeah, so I am a stakeholder in enharmonic spelling. The bug report was from somebody who wasn't a stakeholder in enharmonic spelling. Great. Who might not care about enharmonic spelling? Yeah, go ahead.

AUDIENCE: I would say a guitarist who doesn't know sheet music.

MICHAEL SCOTT ASATO CUTHBERT: Yeah, totally. Or who doesn't-- maybe just be a little more generous. Maybe they know something, but it's translating from the sheet into the tablature into the position before translating to the next stage of mental processing. I mean, we do that a lot. People who are native speakers in two languages don't translate from one language to another before there's some kind of conception. Whereas when I'm learning a new language, it's always going through English before it has a meaning to me.

Yeah, so somebody who plays guitar might be translating what they're seeing to a fret position and then saying, that's great. You see that quite often in a lot of fake book scores, in things where you'll see-- I don't know. We'll say same type of thing.

You'll see E flat minor on something, and then you'll read on that, the chord-- I think that's right-- quite often, or something like that in a melody. And well, obviously, that's not. But because it's all about translating it first. That's good. Good. Other stakeholders or other interesting people who might use pitch in a different way? Yeah?

AUDIENCE: I've heard there's people who make incredible music, but they don't even know music theory at all. It's just like--

MICHAEL SCOTT ASATO CUTHBERT: Yeah.

AUDIENCE: --useless to ...

MICHAEL SCOTT ASATO CUTHBERT: Yeah, totally. There's a big, big area now of study on implied cognitions of music theory. That the idea being for some of the people who are working on this, that nobody who's making absolutely amazing music has no conception of what they're doing. They might not have the vocabulary to share it with us, but there's some sort of underlying conception. And how do we get at that?

One of my colleagues-- one of your fellow professors you can take class with, Leslie Tilley, who teaches Rhythms of the World and the Balinese gamelan and Music of Bali and Indonesia-- thinks she just wrote a book that just won the big, big award in music theory, which is on trying to figure out what the conception of Balinese musicians who play these two people, who play these improvised drum patterns interlocking at huge, super, super fast intervals. But sometimes one of them stops a bit and does the other one, and they almost never miss. And when they do, just like a flubbed note, they laugh and say, oh, how did we screw that up?

And her book begins with asking people, how do you how do you do this? They said, oh, you just play whatever you want. And so of course, you play something. No, no, not that. Not that. And so how do you get at the implied music theory conception? Super great question, Jordan.

I keep saying, who are the people who do this? But some people in this class identified stakeholders that weren't people. Yeah, go ahead.

AUDIENCE: Like DAWs, the digital audio workstations.

MICHAEL SCOTT ASATO CUTHBERT: Yeah. Yeah, so we can give autonomy to a digital audio workstation or to a computer in general, or to a neural network. There's been a lot talk about autonomy and sentience in neural networks this week. And so why not talk about them as stakeholders? What makes it easier for a digital audio workstation to process a pitch? Usually, eliminate an harmonic spelling. It helps quite a bit.

Yeah, great. Anything else that came up that you found? Especially if you read something-- and even you don't remember who it was, but somebody else's that you read and liked and you thought, oh, I wish I had thought of that. I'll see it when we go through the sheets and stuff. So that's great.

That's mostly leaving aside topic one. Now, there's one more thing in topic one, and I've put it on video to watch this weekend. You will have unlocked your first Music21 module. So now in any problem set from now on, it will be fair game to say, from Music21, import pitch. And

The video will go into more detail from this. pitch is lowercase in Python. What's that signified if you're importing science lowercase, Probably Well, actually-- well, OK. There's a pitch dot capital Pitch. What's the distinction between those two in this case? Yeah, Jordan?

AUDIENCE: One is a class.

MICHAEL SCOTT ASATO CUTHBERT: One is a class and the other one is a what?

AUDIENCE: A folder.

MICHAEL SCOTT ASATO CUTHBERT: It acts exactly like a folder. Great answer. There's specific name that that's used in-- yeah, Jason?

AUDIENCE: Module?

MICHAEL SCOTT ASATO CUTHBERT: Yep. So which one's which?

AUDIENCE: That one's the class.

MICHAEL SCOTT ASATO CUTHBERT: That one's a class. I'm very bad at-- I turned around and you're pointing to the--

AUDIENCE: Capital P.

MICHAEL SCOTT ASATO CUTHBERT: Capital P is a class and lowercase one is the module. So you have everything module. Yep. Everything that's in the pitch module is on because you will also find in there things like accidental, which is pretty helpful. So a module groups together related set of classes if you-- yeah?

AUDIENCE: Sorry. Just going back to the representation thing, I was wondering if there are any studies on people who don't have perfect pitch, but only relative pitch. Like, take the same piece of music and then just transcribe it. Is there any emotional difference to the matter of it's like, starting on D or starting on F to someone who doesn't have--

MICHAEL SCOTT ASATO CUTHBERT: Who doesn't have perfect pitch? I'll have to look that up. I imagine-- I mean, I don't have perfect pitch, but absolute pitch is the term that we're using more often now. I don't have perfect or absolute pitch, but there's definitely pieces that you play some song that I know really well, and you play it up a half step or up a whole step, it's going to impact me.

Now, if it's a song, I don't know, I'm not entirely sure. But I'll see-- let me just make a note on that. I found in one unscientific study-- did that for a paper I was working on-- that there's a particular case where you can, very strangely-- the number of notes within a repetition of contour is decreasing but the number of repetitions of that contour before something changes is increasing. So the tempo, depending on if you hear the harmony or when the harmony changes, the harmonic changes as determinant of tempo. Or if you hear the repetition of contour as determinant of tempo, you can hear the same piece speeding up or slowing down.

And a lot of people were like, oh, yeah. I hear this. But in that unscientific experiment, people with absolute pitch said, I don't hear repetition of contour when the notes are different. That doesn't seem, to me, a kind of repetition. So there are probably different ways we can do that. Great. Super great question.

We'll put that aside when we edit it. We'll edit that about 10 seconds back, and we'll jump back to here just for a second. Don't be afraid of doing that. If it's something that I think it's going to be super confusing to everybody if we jump back, I'll say, we'll hold that for a bit. But I think we can still code switch. We can switch back and forth between what we were talking about 30 seconds earlier and what we were just talking about, right? Good.

OK, so you're going to be able to import and use anything in the pitch module, as long as it doesn't have the word transpose in it somewhere because certain kinds of transposition will be things for us to work out in the next couple of problem sets. Good. So that's what's coming up, and then we'll be done with that.

So last class, we were talking a lot more about moving from the representation of individual pitches, individual notes into larger things. We might call it representation of a score or simply representation of music, something of that sort. I introduced a term that some of you already knew from other classes that goes back to philosophy, but has an interesting twist of a meaning in computer science. Anyone remember that term without looking down at the outline?

AUDIENCE: Ontology.

MICHAEL SCOTT ASATO CUTHBERT: Ontology, good. And last class, we explored some of the things in the philosophical definition of essence of what something is. I remember some of the things that we had trouble defining when something ended and became something else. Yeah?

AUDIENCE: We talked about notes.

MICHAEL SCOTT ASATO CUTHBERT: Great. We talked about notes that we can't really tell is this one note? Is this two notes? Is this six notes? Is it an infinity of notes or 14 or-- yeah. Then a larger system of musical notation, remember, we called it used in one part of the world. Yeah?

AUDIENCE: Common Western music notation?

MICHAEL SCOTT ASATO CUTHBERT: Common Western music notation, yeah. And we had some trouble distinguishing when what was common or what was Western about it. Or I suppose in one case, what was music about. It seemed to be on the edge. Good. Then we also brought up a little bit the computer science definition. And that's what we're going to be looking at a little bit now, but always feeling like we want to jump back to the philosophical definition because we don't want ideas from computer science always to-- what is easy to conceive of, we don't want that to limit our imagination of what we might be able to conceive in the future.

So we talked about in the computer science definitions, there are various ones. There's ones that are more formalized. There's ones that are less. One that I keep in my head because I can remember it when I come up is when we're trying to represent something like music-- or you can be your ontology of a website or something. But when you're trying to think of a particular object that's either conceptual or reified-- really real-- I want to think about how is it made up of what objects are in it? What are the relationships between objects? And what are they called or what are their names?

What was it? There was something-- it was like first or second class that we had a real-- oh, I remember now-- really hard time figuring out the name of it. And it had to do with the-- well, this is an eighth note and this one's a whole note. And so the blank of the note is eighth. The blank of the note is whole. What did we come up with? We had some--

AUDIENCE: Almond.

MICHAEL SCOTT ASATO CUTHBERT: The almond of it. Yeah, we had the-- so sometimes pushing it way out there can make that. I mean, some of y'all have taken physics, that probably subatomic particles do not have a charm, a color, things like that. But these kinds of names can help us conceive of some property of something.

So we might have the almond of a note. I think I went with something very generic-- the type of the note. But always avoid type if you can come up with something better because type, it's too valuable to use and too generic. So good, so what are things and what are they called? And I'm spending-- notice, what are the objects? You're going to be thinking about that a lot. What are the relationships? You're always going to be thinking about that a lot.

But the name of something can really, really help somebody else understand how you're conceiving something. And it can be very, very hard. Some of you have probably heard the computer science joke that there's only three really difficult things in programming-- what to name something, and off by one errors.

[LAUGHING]

Good. Good. Some people hadn't heard that, and some other people have. Yeah, so what is something called and that. But we're going to jump now to what are the relationships between objects? What are some things in music that we might consider an object? Let's just rapid fire. Just shout out over each other.

AUDIENCE: Note.

MICHAEL SCOTT ASATO CUTHBERT: Note. Good. Always take the easy one. Get the easy participation credit. Great.

AUDIENCE: Rest.

MICHAEL SCOTT ASATO CUTHBERT: Rest, I heard. Yep.

AUDIENCE: Measure.

MICHAEL SCOTT ASATO CUTHBERT: Measure. Oh, that's good.

AUDIENCE: Phrase.

MICHAEL SCOTT ASATO CUTHBERT: Phrase, ah. Good.

AUDIENCE: Articulation.

MICHAEL SCOTT ASATO CUTHBERT: Articulation. Good. I'm going to stop writing down unless I hear something that I've never, never thought of. Good, what else?

AUDIENCE: Clef.

MICHAEL SCOTT ASATO CUTHBERT: Clef, good.

AUDIENCE: The whole score.

MICHAEL SCOTT ASATO CUTHBERT: The whole score. Good, the score. The dynamic. The instrument. Tempo. Name of the performer. Concert hall it was first performed in. Maybe, maybe not. We can kind of go-- but these are things we might represent when we're thinking about a performance of the music. Good. Anything else jumping on people's minds? Angelica, do you have one that you haven't heard?

AUDIENCE: The word beat or something.

MICHAEL SCOTT ASATO CUTHBERT: Beat. Oh, yeah. Beat. See, that's great because that's something we don't see on the score but we're thinking all the time, right? Yeah?

AUDIENCE: Timbre?

MICHAEL SCOTT ASATO CUTHBERT: Timbre. Good. The instrumental color or timbre. Great. So these are all different things that we might think of as objects. One of the things-- and some of them, we might not represent in every programming language as an actual object class, depending on your programming language.

But there are things we can think of that have a particular noun-ness. Can you say noun-ness? That they act like nouns, that they have the essence of a noun as opposed to a verb. Nobody said that a fundamental object was transpose or perform, those verb-y things.

So what we want to think of various ways that objects can be related. We talk often about taxonomies and meronomies. Has this come up in anybody else's in computer science classes? Doesn't come up in, what was it? 6.009, 6.1010? OK, good. But who has heard of a taxonomy in another concept? Yeah. Yeah, Karima, where have you heard it?

AUDIENCE: Like animals.

MICHAEL SCOTT ASATO CUTHBERT: Animals. Kings play cards on fat green stools. Do they still teach that? Or kingdom, phylum. Yeah. I mean, now there's all these-- they keep rearranging it, right? They're keeping subclasses and super kingdoms, and things like that. But yeah, so we can have a taxonomy that way.

So taxonomy is one that were useful for taxon-- when you're too close, it's hard. OK. Taxonomy, in this definition in ontology, it's something is something else. So give the example here that a minuet is a dance. For those who know the minuet style in the early 18th century, the minuet is a triple time dance, which means it is a dance. Or it is a-- I don't know. What's something else in music that is something else? Yeah?

AUDIENCE: Allegretto is a tempo.

MICHAEL SCOTT ASATO CUTHBERT: Allegretto is a tempo. Good. I put up this. What do you call this one in this context?

AUDIENCE: Staccato.

MICHAEL SCOTT ASATO CUTHBERT: Staccato. Now you'll understand when you do optical music recognition how the context is so important. Otherwise, you would say augmentation dock, right? So staccato is a--

AUDIENCE: Articulation.

MICHAEL SCOTT ASATO CUTHBERT: --articulation, good. Is there any higher level thing that we might say is? Musicians? Probably not. But when we're thinking of representing for computer-- yeah, Sean?

AUDIENCE: A half note is a note.

MICHAEL SCOTT ASATO CUTHBERT: A half note is a note, good. Half note is a note. Great. I was just going to go back to the articulations also, but we'll jump to half note is a note in a second because that's a pretty fundamental one. Sometimes with your own representations, you might think of, for instance, is trill an articulation?

Who votes yes? Who votes no? Who votes I have never thought about this, and I hope I never have to think about this again in my life? OK, good. Good. A lot of people call a trill an ornament, a tern, an ornament, something like that. So you might need what is the class of-- but we know that a trill is more related to an accent mark than it is to a half note or to a Giselle. Right?

So we might think of what's some kind of slightly bigger thing that holds articulations and ornaments? What might we call that? And why don't we come up with a name for that? Yeah, Jason first then--

AUDIENCE: Modifier?

MICHAEL SCOTT ASATO CUTHBERT: A modifier. Good. That's what a lot of people-- believe what it's called in the VexFlow notation software. So somebody's-- yeah?

AUDIENCE: I think maybe you could throw under envelope if you're thinking about synthesis.

MICHAEL SCOTT ASATO CUTHBERT: Envelope. Are they envelopes? They might be envelope changers. They, in themselves, the staccato doesn't-- does the staccato have an envelope? Maybe it does. I'm not-- I don't work in this as much.

AUDIENCE: But it's like, it's just like the attack-- I forget the second thing.

MICHAEL SCOTT ASATO CUTHBERT: Sustained attack? No, attack d--

AUDIENCE: Decay.

MICHAEL SCOTT ASATO CUTHBERT: --decay. S?

AUDIENCE: Sustain.

MICHAEL SCOTT ASATO CUTHBERT: Sustain. R? Resonance or release. Yeah. Not this class. Not what we're talking about, but some people have that knowledge from doing electronic music things. Yeah, so this is something that changes how we attack, how we decay, how we sustain, how we release. So yeah, maybe that is a very good term.

We're going to be coming up with a bunch of these conceptual grouping things because of something that you can do with things that are in a taxonomic relationship. And I'll write that out. And while y'all think of what might be in those missing letters, you had a comment on the previous one.

AUDIENCE: I just want to say augmentation.

MICHAEL SCOTT ASATO CUTHBERT: They might be an augmentation, or in the case of a staccato, maybe a negative augmentation or something. Yeah, a length reducer. Good. What do you think that in computer speak, when we say that something is something else, we can represent it with subclass relationships. Subclasses is something that comes up, I'm 99% sure, in the prereq for this. Right? You guys deal with classes and subclasses and things.

So quite often, there are a bunch of built in subclasses in Python that you have various-- a lot of things that we think of, things are subclasses of number, complex numbers. Oh, yeah. We all learned taxonomies in our math classes too, right? Complex numbers, floats, ints, things like that.

And quite often, for different purposes-- that's why we thought about stakeholders. For different purposes, different taxonomies can be created, right? I mean, in certain cases, you could think, well, an integer is a special kind of float that has nothing past the dot. And that's a very great mathematical way of thinking of an integer.

If you try to implement it that way in your computer programs, you got some slowness going because you have a whole bunch of added baggage that you wouldn't need. Or you could think of if you are a-- I don't know. When you're in third grade, fourth grade, when's the first time you learn about decimals? We don't know the word int, but we know that there are integers. We know that there's 5 and 7 and things.

And the first time that you see 5.2, you would think, oh, well, this seems to be a subclass of this because it's a different kind of number. So depending on how our ontology, how we're conceptualizing the world, taxonomies could run in two different directions. So we'll get back to this. I just want to make sure that we can get through this in a certain amount of time.

Then there's something called-- where did I write it down? Did I not write it yet? No. Something called a meronymy. And a meronymy, something that instead of is, it's something that something has.

Michael Cuthbert is a human most of the time, and therefore, he is an animal because all humans are animals. Therefore, he is a living thing. Or that is because we can think of that as a taxonomy. Michael is an arm. He's actually two arms, and he's also two legs. No, that doesn't make sense. There is a relationship between my arms and my legs and me, but it's not the is relationship, or at least, not the way I conceive of it. But it's something that I have.

So what is something-- oh, sorry. I meant to leave up John's great half note is a note thing. So a half note is a note. What is something that a half note has that it is not? Jason first.

AUDIENCE: Almond.

MICHAEL SCOTT ASATO CUTHBERT: Almond. It has almond, which we used for the type. Right? That's the almond. Great. I'm going to-- I already have chives as a vocabulary later, now we have almonds. So it has a particular type. Great. What else does a half note have? Adam?

AUDIENCE: Duration.

MICHAEL SCOTT ASATO CUTHBERT: Duration. Good. Might have a duration. What does an eighth note have that a half note lacks?

AUDIENCE: Flag.

MICHAEL SCOTT ASATO CUTHBERT: Flag. Call that a flag. And what might two eighth notes have that one eighth note doesn't. Yeah?

AUDIENCE: A bar.

MICHAEL SCOTT ASATO CUTHBERT: A bar, or what do we call that thing?

AUDIENCE: A beam.

MICHAEL SCOTT ASATO CUTHBERT: Beam, good. Good. I'm going to-- the center of the room has been really active for a bit. I'm going to focus over here and over here for a little bit for the next couple answers. Good. So it has a beam. So we can think of this.

And here's where your conception skill-- eighth note has a beam. 16th note have a beam, or 16th notes have a beam? Yeah. Do 16th notes have a beam that has two lines in it, or do 16th notes have two beams?

[LAUGHING]

You can imagine you're going to encode this very differently if you encode the number of things. And some eighth notes don't have beams, they have flags instead. But they have the potential for gaining a beam, perhaps, if we think of-- mommy, I always want to grow up and be a beamed eighth note, not a flag eighth note.

So you might say, well, we're not going to have a certain eighth notes that have no-- now we'll get into computer speak rather than philosophy speak-- that have no attribute that where beam could go because that's just silly. I don't want to have unbeamed eighth notes and beamed eighth notes. Maybe that's two different concepts. That's two different objects.

Maybe you say, well, I'll have one type of eighth note and it either has or doesn't have beams. Then in your representation, you might be asking yourself-- oops. Leave that up. You might be asking yourself, does it make sense to say that eighth notes, 16th notes, 32nd notes, 64th notes, whatever-- go on to the highest power of two you can think of- have beams, but quarter notes and half notes and whole notes don't? Or do we say that whole notes have the potential to have beams, but you better not put any on?

I don't know. There isn't a correct answer for that. And so sometimes, what we conceive of, we-- I cannot conceive of a whole note with beams. I can conceive of a half note with beams. I can think of it. I don't know what it means. Actually, it means [INAUDIBLE] tremolos. Dee dol, dee dol, dee dol. But that's something else. Maybe that's something visually that looks just like a beam, but means something-- or maybe it's not.

So some of the things that come up in musical forms give-- so a sonata is a type of musical form. It's used quite often in the 18th and 19th century, and it's a musical form that has an exposition-- the thing close to the beginning. And the exposition has a first theme group and a second theme group. And the first theme group has a phrase, and that phrase has measures, and those measures have notes, perhaps is one way of conceiving of it.

It would make much less sense to say, well, it has a phrase. And a phrase is a measure, and a measure is a note. Right? So this is going to be the big part of problem set two. In fact, we're going to skip some other things. We're just going to do a little bit of the cook discussion. And then I want to give you time to start thinking about problem set two.

In problem set two, you're going to think of a type of music that is not common Western music notation. It might not even have to be a type of music. Some people have done things related to music, such as there was a chess playing music generating system in the 1960s that somebody decided to work on. Or it might be the structure of a concert or something like that.

And I want you to come up with the kinds of what's the taxonomy in the ontology in the conception of the system? How does this work, and what are these relationships? What's the meronomy and the taxonomy there?

Over the weekend, we're going to be going through some of the past notational formats for common Western music notation, some of the representations, the ontologies, the file formats that have come up over time. One of the things-- a lot of the pieces, the discussions that people were talking about have some connection to common Western music notation, things I was hearing, but are not themselves common Western music notation.

They might have-- oh, yeah. Well, the rhythm is the same as common Western music notation, but something else. Anything that you think is just part of common Western music notation, assume that part of the ontology of the file format has already been solved.

So what I'm saying, if it's like you're doing something with, oh, this is a dance groove. And the baseline is stored this way, and it's represented in this way. And the notes of the baseline are just stored in CWMN, or something, or stored somehow-- stored in MusicXML, something. You'll get some formats.

The only way that people have gone really wrong on similar assignments to this in the past is spending too much time thinking about the things that have already been solved. If your format needs to encode in images at certain points, just say they're JPEGs. Don't invent your entire new image format. Don't invent new things that you already know that have-- spend the time on conceptualizing what hasn't been conceived before.

I just want to point out a little bit-- oh, sorry. That is completely illegible. Always feel free to tell me when something is completely illegible. Is that-- there we go. Now it's not. We're going to go through this more slowly. I think I can't remember over the weekend or on Wednesday. I should have my syllabus in my head, but I write it down so I don't have to memorize it.

But we're going to be going through something called the Rosetta Stone of Digital Music Representations by a really, really smart guy, Craig Sapp, who realized-- anybody know what the Rosetta Stone is? It was the same text written in three different languages, especially in Greek and in ancient hieroglyphics, that allowed hieroglyphics to finally be decoded. Craig Sapp's Rosetta Stone takes the exact same little melody. We can go over-- what is it?

[PIANO KEYS PLAYING]

My rhythm was pretty bad on that. But that exact same melody, and shows how that melody is encoded in-- I think he has 32 on his website. I have edited it down to about eight so that you can see what different representational choices people have made, and what kinds of ontologies they make.

For instance, I remember in this case, beams are high level things. They are represented by these braces. And so therefore, they enclose things. Beams, in this case, seem to have a meronymy that beams have notes, not notes have beams. So we'll try to see some of these things.

So another format called ABC-- one of the ones still in use, even though quite a bit old because it's very, very simple. We'll be looking at these in more detail, various formats that at some point were very big, but have mostly faded. This is one that we saw once before, MuseData, where the Q was quarter and the eight was a certain number of durations, and F5-- do you remember that from a class or two ago? Yep.

And we'll see over time, how the formats have gotten to be a little bit more complex. Here we go. This is MusicXML-- which we're going to be talking about quite a bit over the course-- invented by an alum of MIT, Michael Good. And one of the great formats for representing music notation, the first one with enharmonics that just was taken and used by almost everybody. It's about 20, 22 years old.

A lot of us thought, this is the final format. This is the representation that really, really works. Even if there's some weird quirks, we use it forever. And then a terrible thing started happening to music. A great and terrible thing started happening to music about 10 years ago, 12-- almost 15 years ago. But we started looking at sheet music on various sized screens. And therefore, where the screen is too small to look at a whole page at the time. Or you turn it sideways, and suddenly your page has changed.

So the notion that music cannot just be represented by where it is on a page, but has to be able to keep continuously reflowing, has caused a big change to people today trying to come up with what is the optimal format for representing music whose presentation might be different for you because you have an iPad, and it might be different over here because you have a 15 inch monitor, and it might be different for somebody else because you've just printed it out in the old fashioned way? How do we represent music that way?

I wanted to point that out because there was a huge debate. So there's a big discussion happening right now on what the music notation next format will be, and exactly how it will-- what its ontology and meronymy will be so that it survives into the next generation of musical scores. And people, you can see, are still right now-- some idiots sometimes post. And exactly what kind of format, what the advantages in terms of space, speed, various representations are going on. And I haven't checked it this morning, but it seems to just still be going on.

So this is a continuous topic of discussion. People are still very interested in what's happening, what music representations are out there. And I'm looking forward to five new music representations existing a week from now that don't today. Thanks, everybody.