CMPS 10 Lecture Notes: Lecture 6 (1-21-2016)

Clicker Registration Fun

Did you guys get the clicker announcement?

- Yeah

Teacher has never used iClicker thing before, so hopefully it goes okay!

- First let's see if we can get everyone to register.
- There seems to be some difficulty.
- For people with iClicker+ shutting it off and turning it on again seems to help.
  - Also, if you tried to register and failed, pushing DD twice will make you unregister so you can try again.
- There is something strangely hypnotic about watching the names light up.

Last Time

- Made the claim that if we have a continuous function of time measuring the pressure on our ear drum, even though that function is continuous, if we can measure it at evenly spaced intervals, so that instead of a continuous plot you only get to know certain dots (i.e. the x and y coordinate), you can capture the function perfectly.
  - And because we are measuring regularly, really we only need to know the y values.
- And we say that we can measure the y value with infinite accuracy.
- So let's say that we've measured our function at 100 points in time.
- Teacher would like to make the claim that, if this function is representative of sound, and teacher wants to capture the function with perfect accuracy from point of view of human fidelity.
  - Then we would need to sample the function 40,000 times per second (i.e., 40,000 points per second).
    - 44.1 KHz is roughly the same.
- **The deeper claim:**
  - Though the continuous function is an infinite amount of information, if we sample it densely enough, we can perfectly reconstruct it, without any accuracy loss.
  - Can we think of a situation that we are familiar with, in which we are dealing with a continuous plot, but a finite number of points if sufficient to capture it?
    - Let's say, for example, that we are promised that the function is a straight line. How many points do we need to describe it?
      - Only 2.
    - Given two points (and we only even need the x coordinates of them!) Once we have that we are set!
      - Even though it was a continuous function, we could describe it perfectly with only two points!
    - What if it was a parabola?
      - We would need 3 points!
    - If you look at the function of the straight line: y = ax + b, we needed two points.
      - i.e. if we knew a and b, we could compute y for ourselves.
    - IF you look at the function of the parabola: y = ax^2 + bx + c we need three points.
      - i.e. if we get told a, b, and c, then we can compute y for ourselves.
    - So we have these lines that seemingly require an infinite amount of information to describe, but we can actually communicate it perfectly with a very small amount of information!
  - But keep in mind: the above really only works if we are OK with ONLY drawing a straight line, or ONLY drawing a parabola, which isn't really how it works in music.
- But in the case of music, we are interested in capturing everything; bird song, orchestra, human voice, etc.
  - It's a much, much, bigger (infinite) set of possibilities.

    - Also: we don't want to get freaked out by infinity, because even with just lines and parabolas there are an infinite amount of those too. So infinity in and of itself isn't the hard part.
  - So how can we do it? We *need* to have some fundamental assumption about the nature of the function.
    - There is a correspondence between how complicated the structure of the thing we are trying to capture and how much information we need in order to do it.
      - So we cannot be arbitrarily complicated. There needs to be a bound.
    - Any idea how to create a barrier to limit the complexity?

On the Subject of Eardrums

We said that we had an eardrum, and that the ear drum is moving up and down, and what we are plotting when we draw our musical function graphs is the distance from your ear drum's resting position as induced by sound.

- Do we all agree that our ear drum has mass?
- Do we agree that that puts an upper bound on how fast it can possibly move?
  - In particular: how fast it can possibly change directions.
- Let's go to basic physics:
  - If no force is acting on something, what is the behavior of that thing in space
    - It is either standing still, or it is moving at constant speed.
      - Example: Curling: if there was 0 friction, then the thing would move at a perfect speed and there would be no acceleration or deceleration.
- So let's say your ear drum is going up
  - It's speed is a vector going in one direction.
  - And once it reaches as far as it can go, the speed becomes 0
  - And then it starts moving in the opposite direction.
  - It means there is force applied to your ear drum.
- So if your eardrum could move super fast, it would require an infinite amount of force, which would cause your ear drum to burst.
  - A loud sound causes your eardrums to move so fast they burst, which is what causes you to go deaf.
- So there is a band of acceleration in which your ear drum can move
  - There is a bound on how fast it can move before it tears.
- In mathematical terms: what does it mean to say that your eardrum can't move infinitely fast.
  - Y can't be infinitely high? WELL, sure, but that wasn't what teacher was asking about
  - Teacher wanted to know that SLOPE of the curve can't get TOO vertical, because the slope is the speed of your ear moving
    - Big slope means moving quickly in a small period.
      - If y axis is movement in space, and X axis is movement in time, then the slope is how fast your ear drum moved.
    - Smaller slope means your ear drum is moving slower.
      - So slope of the curve is the speed of your eardrum.
    - How many have seen that speed is the derivative of distance with respect to time? Well, now that makes sense!
  - So claim is this: because ear can only survive so much force, there is a maximum speed it can handle, and because there is a maximum speed it can handle, the slope of the curve is bounded.
    - And THAT is the assumption from which we will extract our claim
    - That if we measure the function sufficiently densely in time, we have frozen it, determined it uniquely, without seeing all the in between points.
  - Essentially claim amounts to the following:
    - BECAUSE it's music, it means the slope can never exceed a certain angle.

    - And therefore, if we measure it, where the density of the measurements depends on the angle (the greater the angle, the more densely we need)
    - Then as long we do that, there can be no funny business
        - Have to stay within a particular bound. And if it were to go outside the bound would require it to be very 'agile' which would violate the assumption that the speed is limited.
        - Function HAS to have simple behavior between two points, because to not have simple behavior would require the derivative of the function to become very large.
- So if we sample densely enough, we can freeze it perfectly in time, without knowing every point of the function.

Multiple Choice Clicker Questions!

- If someone tells me an upper bound on how rapidly my function can change in time, and from this upper bound I can extract density in time which is sufficiently big, so that if I measure my function at this density, then I literally trap my function and determine it uniquely, even though I don't know the value of the function in the in between points, I can still figure it out.
- So first claim: based on information presented so far, does this claim appear plausible
    - A for Yes, B for NO.
    - Roughly 83 percent think YES, and 15 percent think NO.
        - Also note that we are currently in AP mode, which means anonymous polling mode.
- Can someone who says NO say what the stumbling block is
    - There is no proof going on right now, but there are two ingredients:
        - Ingredient 1: There is a bound on slope (it will always be as slow moving as THIS or slower)
        - Ingredient 2: we take more and more samples evenly spaced, we claim that based on the fact that the slope is bounded, but there exists a function where the most naive solution to connect the dots, IS the function. If the dots are sufficiently close (where definition of sufficiently is based on what the value of the slope is).
- Teacher can not over estimate the importance of this
    - There are so many things in the world that we view as continuous.
        - But if we have a finite computer with only so many things, we can't represent it right?
            - No, that is BS. World is not continuous, it just seems that way because we are stupid. World is discrete.
- We think our brain works on 1,000s of Hz instead of the Millions of the computer
    - The difference is maybe that human brain has a ton of CPUs in it.
- Let's take human pregnancy for an example
    - How many neurons are created in a baby's brain in the course of a pregnancy per second.
        - How many people think its more than 1 per second?
    - 10 per second? 74 percent say yes
    - 100 per second? 70 percent say yes
    - 1,000 per second? every second that goes by, 1,000 fresh new neurons. Day in, day out, for 9 months: 23
        - Answer: 7,000 neurons per second! Those are processing units that are being in the baby!
            - And note: most neurons ever created happen during pregnancy.
- We draw attention to this because there are bias against what the computer can represent since we think the world is continuous
    - But no, we can represent anything! And due to things like what we saw before, we learn anything.

What about Dog MP3s? How would we do that?

- We could do it just like above, computing the new slope for them. Would it be bigger or lower?
    - Ultrasonic hearing is what Dog has, essentially just means better hearing

  - So slope of dog's hearing is better than ours.
  - So if we had a microphone that was able to record sound as good as our dog could hear, and then if we only sampled it at the rate that is good enough for humans, your dog would be unhappy with the sound quality.
    * To make your dog happy, you would have to sample at twice the rate. Because they can hear up to 40 KHz.

How to Actually Measure That Y Value

Now, remember, earlier we talked about being able to measure the Y exactly, but in general we can't measure the Y exactly (Remember: Y is position of ear drum).

- So when you measure your function, and you've placed a dot on it, there is some amount of precision you can afford.
  - And the way you measure it is just like with those same tree structures we had before.
- So your graph goes from +1 to 1.
  - And in our tree, left means "down" and right means "up" (i.e. lower half or upper half)
    * So first question is, upper half or lower half (i.e. positive or negative)
    * Then next question is, it is in the upper half or lower half OF that super half.
      · And so on and so forth.
- Just like before, adding more children to the trees will allow us to get even more accuracy.
  - And voila! We are back to what we had before with the LEFT/RIGHT decision structure.
    * If we have 16 'bits' then we take the number from 1 to 1, we chop them up into 2^16 (or 65, 536) intervals.
      · And so 16 bits teaches us one DOT, one sample
      · And with 40,000 of those, then we get ONE SECOND of audio.

New POLL: How comfortable are you with the above claim?

- There are a lot of people uncomfortable with this claim!
- So we will repeat the claim.

We have a function, x axis is time, y axis will be between 1 and 1 (or, 0 1 does not really matter).

- Y function is position of ear drum.
- We want it to be as accurate as possible.
- So let's chop the y axis into equally sized intervals.
  - So to measure means to figure out in which bucket does the actual value fall.
  - And now, instead of thinking of chopping the range into intersections from left to right, instead do it as a bunch of bisections
    * i.e. cut things into two.
    * Idea of recursion: do it again!
- So imagine that we're trying to record some numeric quantity that lands between 0 and 1
  - and we label each 'bucket' each intervals as a number (a, b, c, d, etc.)
    * And we can imagine that if we had another layer of sub intervals, we could divide the letters into a1 a2, b1 b2, etc.
      · Which would afford us more precision.
- And now imagine that we are going to play our game of 20 questions (or, in this in class example, just four questions)
  - And each question is just like earlier in lecture: are we moving to the left or to the right.
- The amount of precision that we have is based on 2^questions (and in music, 2^16).
  - If we go 16 questions in, then there is no point in refining it further. Humans cannot perceive anything more.

- So 44.1 KHz at 16 bits. This is CD quality.
  - So we believe that humans cannot hear above 20 KHz. Why do we sample at 44.1 KHz.
    - Imagine that there is a playground with a swing, and you can push the swing.
      - You cannot make the back and forth movement arbitrarily fast.
      - Similarly, your ear cannot accept energy higher than a certain rate in time.
      - there isn't an acceptable form for that energy.
      - Just like how if you push the swing very fast, most of the time you would just be pushing air, only rarely would the swing be there.
      - So there is a theory that if you can only accept something at X frequency, then sampling at 2X frequency is what you need to do.
- So if you only sampled at X frequency, would it sound like garbage?
  - Depends on the content!
    - Measuring just a drum would be fine, but if we measured a top hat, then it would be a problem.
  - Just because 20 KHz represents the limit of human hearing doesn't mean that everything we say uses all of that.
    - Phone throws away everything above a certain frequency.
      - That is why sometimes on the phone you hear NOTHING above certain frequency, because they only care about speech.
      - phone samples at significantly less than CD quality. It doesn't care.
      - It is allowing for

How many people have heard about Super Audio CDs?

- Not much. We are a ruined generation. Spend some money on nice stereo equipment.
- So everything that we are saying is that it is heard and stored and trasmitted later.
  - But to say that things are perfectly trasmitted is crazy. Bits are lost all the time, say, on your cell phone.
    - Yet communication doesn't break!
  - So the question is this: Can we achieve reliable communication over unreliable channels.

New Topic: Redundancy

Let's say that we have a lot of groups of 16 bits.

- and a bunch of these groups concatenated with each other.
- Before we were saying that we could ALWAYS see it.
- But what if sometimes we mis hear? What if he shouts it in a noisy environment?
- What could be done or said?
  - He could say everything twice or 10 times!
  - It wouldn't completely eliminate the possibility of a mistake, but it would make it significantly lesser!
  - If we were to use the 10 time protocol, what would you do?
    - You would hear a sequence of symbols, and sometimes, because of noise, you might hear RRRLRRRRR
      - That means that there are two possibilities: either we were trying to say RIGHT and the LEFT was a mistake, OR that LEFT was correct, and ALL the rights were a mistake!
      - But clearly left being the mistake is more likely.
- So in order to achieve good communication, we have to play the odds. There will always be some uncertainty.
  - Let's say, every time teacher says something, Right or Left, there is a demon that picks it up, and with probability 1 in 10 flips it over.
    - So with probability 1 in 10 you hear the opposite of what was said, 9/10 you hear the correct thing.
      - So we received 100 bits, 90 of them would be right, 10 of them would be wrong, and we don't know which ones.
      - And of course, we DID know which ones were wrong, then we'd be done. You'd just flip the wrong ones and you'd be done. But we don't know. We have no clue.

Next Time:

- Insead of sending 100 bits, in order to communicate 100 bits, teacher will send us more than 100 bits.
  - IF we get a certain amount of bits, and no more than BLAH are corrupted, we will learn how much extra information to send.
- Promise for next time:
  - teacher will send us bits, say 1000, but we don't know in advance what they will say (because if we did there wold be no point in communicating)
  - So before he does that, there will be a communication code. And according to this code, teacher will send us not a 1,000 bits but, say, 2,000 bits.
    - And the promise is that the bad guy sitting between sender and receiver, is not allowed to flip MORE than 200 bits. They can flip between 0 and 200 bits.
      - And the claim is that even though we have no knowledge about what bits the bad guy flipped, we will be able to receive our message with perfect accuracy.

Private tutoring

OTSS Online Tutoring Signup System

- Private tutor for the class is named Justin.
- Justin did fancy kaleidoscope last time.
- Student ID, birth date birth year birth month
- allotted an hour either every week or every other week.
- Dates that he is allotting for are tuesday thursday and fridays
- JBTSE@UCSC.EDU
- Times vary.

One last question:

- This class is too mathematical for me

A: YES (22 percent) 35 B: NO (76 percent) 10
Second question: only people for A
A: I wish it was less mathematical (I want change): 3 27 percent
B: I am happy with out it is (it's okay as it is): 8 73 percent

Next question: sick and tired of music stuff I want us to talk about something else
A: I am sick and tired of music stuff: 31 percent
B: I would like to talk about something else. 62 percent

How people know what Shazam is?

- I would like to know how Shazam Works is A: 62 percent
- I don't care is B: 33 percent

Follow up, among the people who voted B, an honest assessment:

- There was a moment in time when we learned what Shazam was. If prior to learning about it, if someone told us that there is an app where you can lift up your phone and it can understand any song in the world from just listening to it for 10 seconds. So among those people, how many of you would have believed that claim vs. no way.

A: I think that should be possible maybe 43 percent
B: No Way Maybe 57 percent

Can someone who felt that Shazam be possible be able to say why?

- Answer: well, computers can do a lot of s***

  - Well, what about Captchas (slightly distorted characters). The ability to do that is supposed to be uniquely human. But the ability to figure out any song in a weird train station, a computer is supposed to be able to do that.
  - Of course, the state of affairs is that Captchas work, and we DO have Shazaam.
    * But there is no human that can do what Shazaam does, but any human can do the captcha.
  - On the surface, the two tasks aren't that different.
    * But one is so easy for humans and the other impossible for them, and vice versa for computers.
    * So this shift in understanding what computers can do is actually very recent.
      · And one of the things the teacher is hoping, is to develop a better intuition about what is easy for computers vs what is hard for them.
      · And the reason why this is far more valuable than understanding HOW computers work, is because many of us end up becoming NOT computer scientists, but it would be great if whatever domain we go into, we understand "oh, a computer can do this better than me.

- For hundreds of years, biologists were designing experiments so that the outcome of the experiment would be easy for a human to interpret.
  - eg gene sequencing.
  - Take many copies of the same string, read it by touching it, as you touch it, it breaks.
  - To Be Continued!