Slide 1

The impact of spectral resolution on listening effort revealed by pupil dilation

Matthew Winn
Jan Edwards Ruth Litovsky
University of Wisconsin-Madison

Acoustical Society of America Fall 2013

Waisman Center

UNIVERSITY OF WISCONSIN
SCHOOL OF MEDICINE AND PUBLIC HEALTH
Department of Surgery

Slide 2

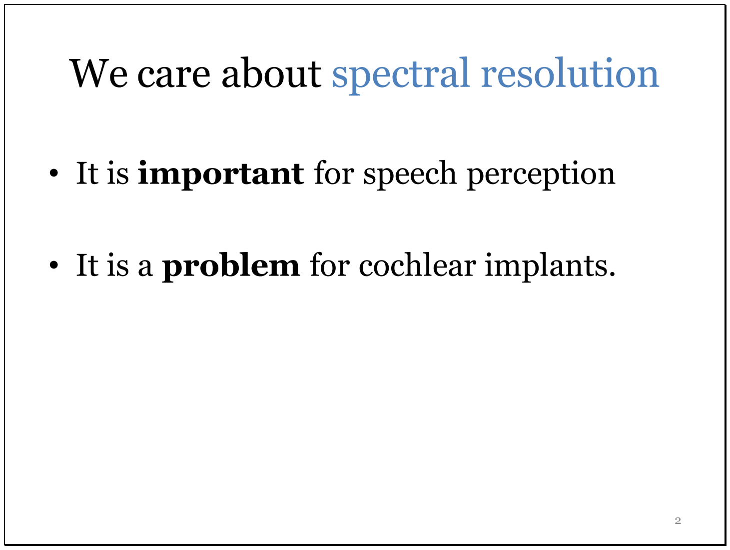

We have a special interest in spectral resolution because it is very important for speech perception – frequency contrasts let us hear the difference between consonants like ba, da and ga, and all the vowels.
This is an especially important topic for me because spectral resolution happens to be the greatest limitation of cochlear implants.
I want you to keep this in mind as I walk you through this presentation today – that we're interested in the impact of poor spectral resolution.

Slide 3

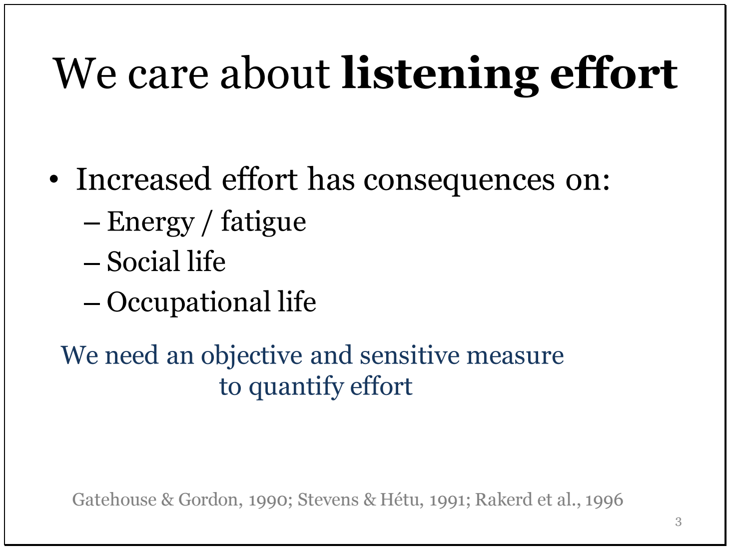

Specifically, as you can tell from the title of the talk, We are interested in listening effort. We know that it takes more effort to understand speech when you have hearing loss. But this isn't just restricted to an experiment in a sound booth. It's a constant, chronic problem, and people who have to deal with this – people who have to focus and concentrate just to hear conversation, it can be exhausting.
People with hearing loss need more time to recover after work, they take more sick days, and they tend to not socialize as much as their normal-hearing peers.

And think about it – for every single conversation, you need to focus, to concentrate, or you'll miss a word and need to guess and fill in the gaps as you go.
At the end of the day, you might have nothing left. And I've seen this – I've had clients who, when I ask "so what do you like to do after work, on the weekends?" They say ehh nothing, I like to jut sit at home, alone, watching tv, you know. I used to like going out with friends, going to the theater, but now I just don't have the energy. I don't want to sit there and miss the conversation, miss the punchline of the joke – That's embarrassing, and it's frustrating.

And it's frustrating for me. It's one thing to measure how many words someone gets right on a hearing test, and it's another thing to have an index of how hard they have to try to get that score.

THAT is what this research is all about.

Slide 4

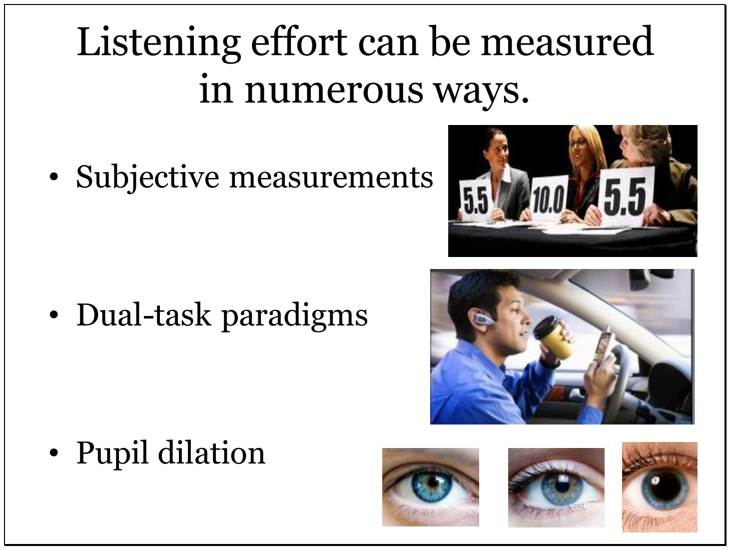

There are numerous ways that people use to measure listening effort.
Traditionally, you just ask… or have people rate on some kind of scale. This has obvious pitfalls of subjectivity and personality differences.
You can also see how effort exerted to complete one task interferes with your ability to perform a secondary task. This is becoming more popular lately.
The method that I have grown to use, and which will drive the rest of my talk today, is measuring pupil dilation.
Simply, when people exert more effort, their pupils dilate more.
Now, you may already know that pupils also react to different amounts of lighting, and whether someone is excited, but, as you'd do in any experiment, if you keep good control over those other external factors, you can get very nice reliable data on cognitive effort.

Slide 5

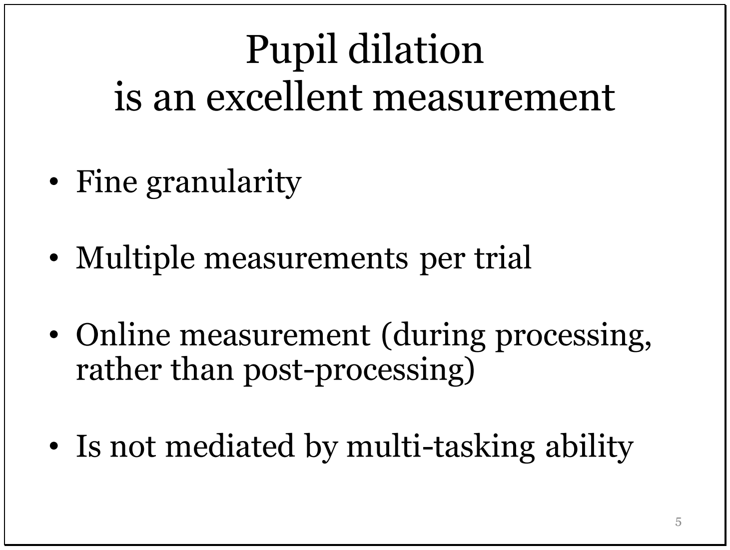

Pupil dilation is a fantastic measurement because it offers extremely fine granularity – you don’t have to rely on a person getting an answer completely correct or incorrect because it is a continuous measure.

You get multiple data points per trial that you can analyze in a time series, and therefore can look at what’s going on *during* processing, rather than measuring the result of processing.

Finally, in contrast to the dual-task paradigms, you aren’t measuring something that’s inherently mediated by multi-tasking ability, where the modularity of different skills can be difficult to unpack.

There are challenges in collecting pupil dilation data, but I’ll show you that they are worth the trouble.

Slide 6

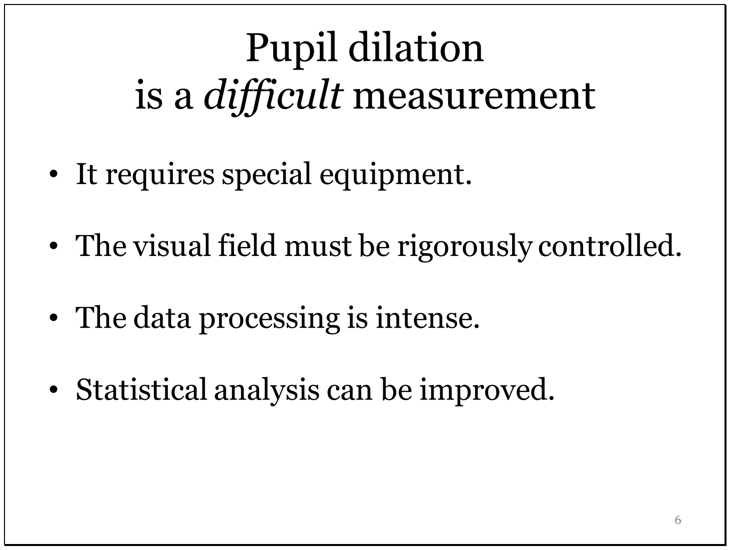

Pupillometry isn't an easy measurement to collect. It requires specialized equipment, some pretty intense data processing on large datasets,
And rigorous control over the visual field, which is something we auditory scientists rarely have to think about.
Statistical analysis of pupil dilation has been undertaken without a whole lot of sophistication, and only recently have we entered a new era of popularizing time-series hierarchical growth curve analysis, which is really the most appropriate way to look at the data.
So, there are these limitations, but I'll show you how pupillometry has revealed some very interesting findings with regard to spectral resolution.

Slide 7

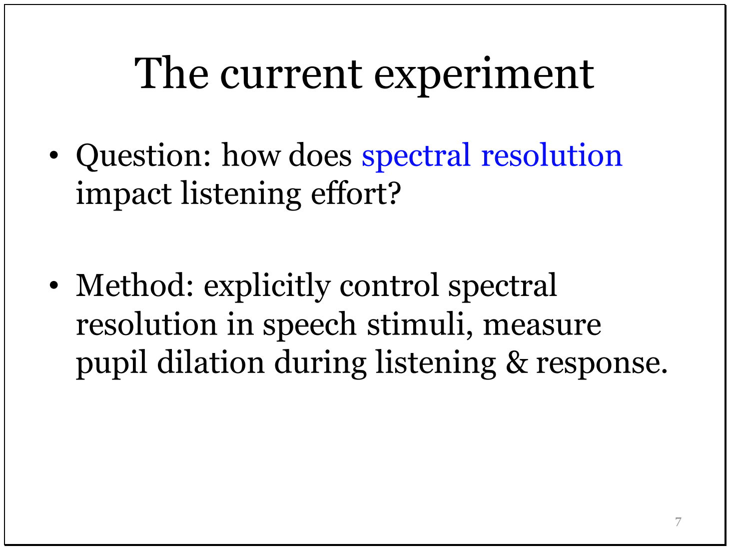

I'll remind you that we're interested in this big problem of spectral resolution, and what effects that has on listening effort.

We controlled spectral resolution in two different ways, which I'll describe in a moment. The idea is, as we systematically change spectral resolution, we measure pupil dilation and perhaps see a corresponding systematic change.

Slide 8

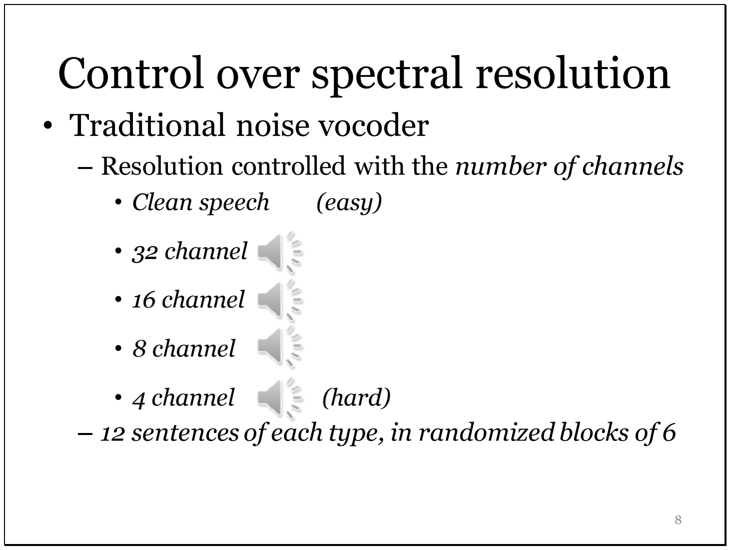

The first way that we controlled resolution was by processing the speech with a noise vocoder.
The vocoder takes the signal, and divides it up in the frequency domain into a set number of channels – you can think of this sort of like the amount of pixels on your tv screen. The more you have, the clearer the resolution.
So we present some normal sentences and these vocoded sentences and measure pupil dilation during listening.

The listeners were 20 undergraduate students with no experience listening to vocoded speech although we do provide a few minutes of training before the experiment begins.
The listeners heard counterbalanced blocks of sentences that were grouped by spectral resolution, so in other words, a block of 16-channel speech, a block of clean speech, a block of 4-channel speech, a block of 32-channel, and so on.

Slide 9

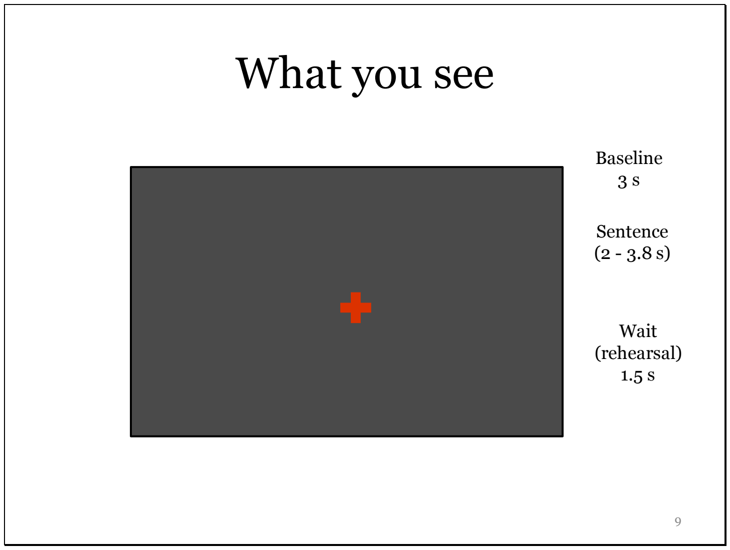

During the trial this is what you see – it's not very visually exciting, because that would cause unwanted changes in pupil response.

First, 3 seconds of silence, where we measure baseline pupil size…
You hear the stimulus, which is a sentence from the IEEE corpus. …
You wait a second and a half….

Slide 10

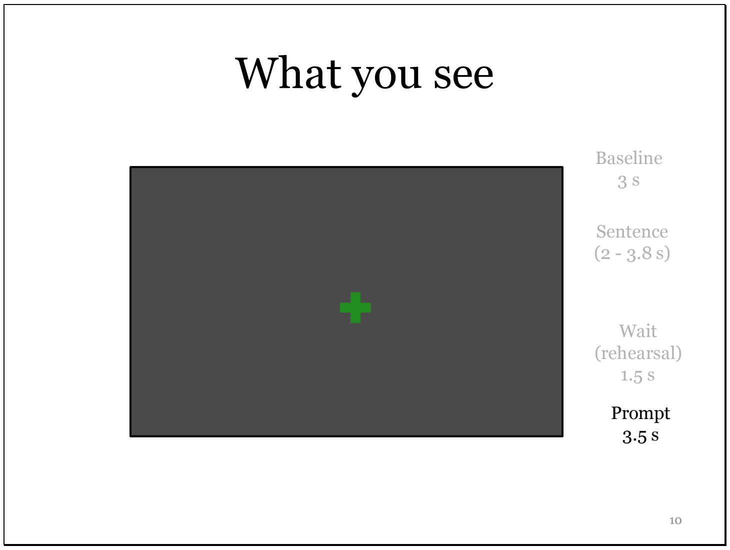

And then you are prompted to repeat the sentence back.

Slide 11

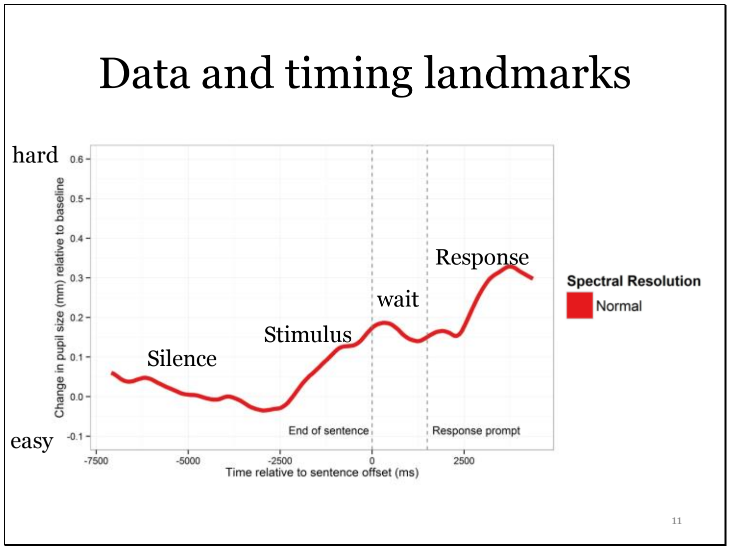

This is what the data look like.
Again we're seeing the data unfold over time, reading left to right. First we have the silence, then the stimulus, the wait, then the response.
This is the pupil dilation for clean, normal sentences. No vocoding, just typical listening.
Note that there's a zero, and that doesn't mean the pupil is not there – this is a measure of change over time relative to the start of the trial, where a baseline measurement is taken.

Slide 12

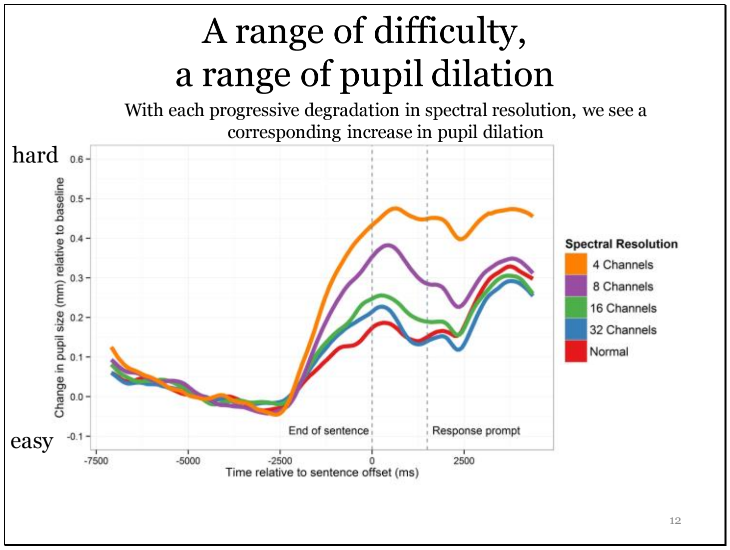

With each progressive degradation in spectral resolution, we see a corresponding increase in pupil dilation

Slide 13

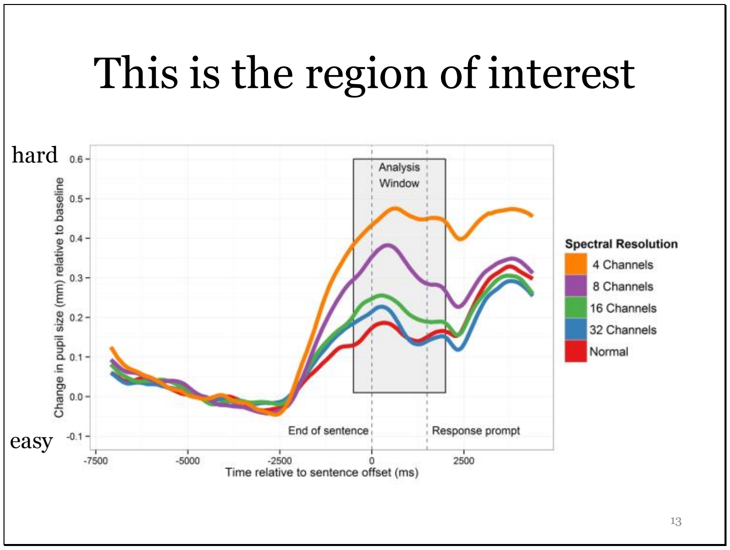

Traditionally, we look at this narrow window, which is the pause between stimulus and response, where the maximum pupil dilation is usually observed.

Slide 14

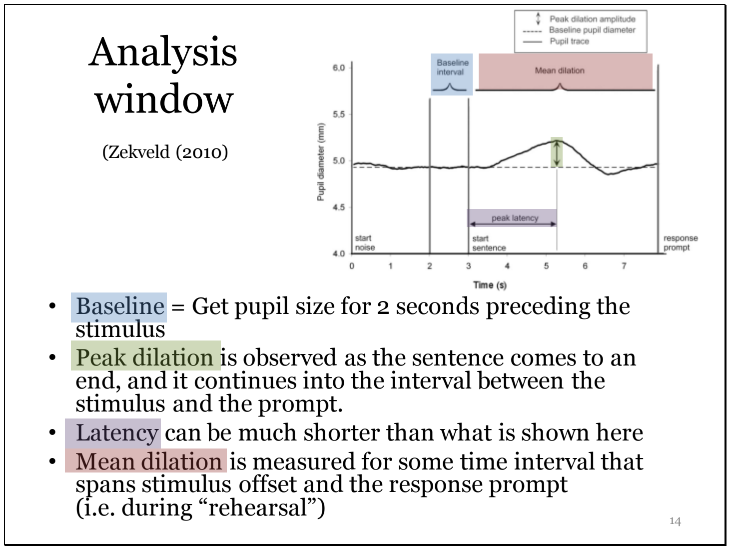

Here's what you find within this window, and different kinds of measurements that are relevant.

Slide 15

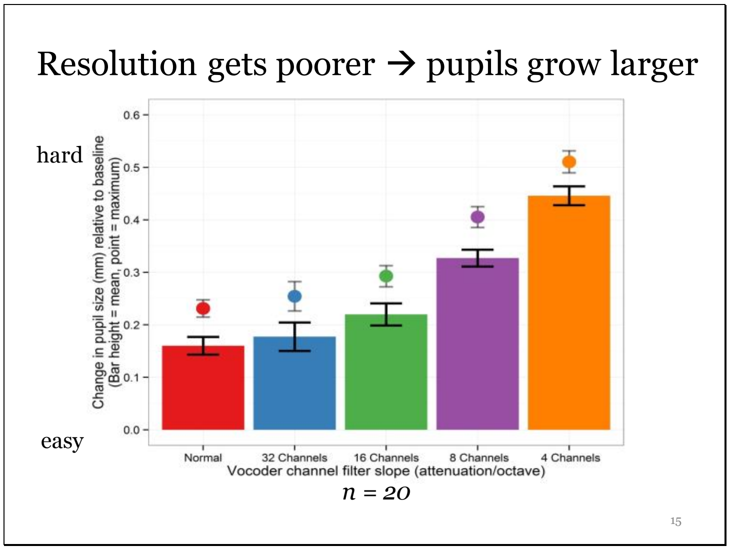

We take the mean value within this analysis window and just re-arrange those values so we can see more easily side-by-side.
... with error bars to represent variability across subjects. The single points above each bar reflect the maximum of each response, and each of those points has an error bar as well.
Historically, both the mean and the max have been used to compare conditions, with the mean being understandably more stable.
We have some reliable differences here, but I'll wait just a few minutes before I get into the statistical analysis.

Slide 16

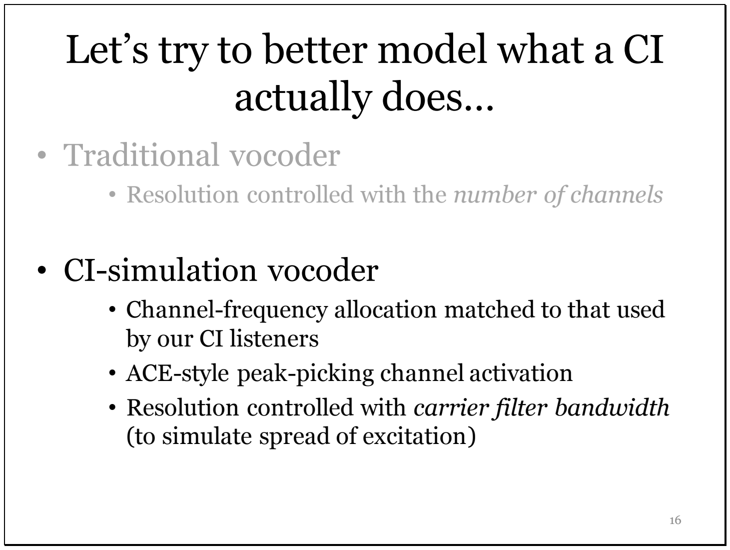

If you recall, I said that we used this *traditional* vocoder, where we altered the number of channels.

The literature tells us that if you use a vocoder with 4-8 channels, you'll get performance that's similar to that of a cochlear implant user.
In reality though, cochlear implants don't just deliver 4-8 channels. They work much differently.
So we wanted to use a vocoder that actually simulates some of the front-end processing of a cochlear implant.
In our lab we primarily work with people who use the Cochlear device, so we stuck with that as a model. We matched the input frequency range, matched the style of channel peak-picking, and, since that was constant, we changed spectral resolution by simulating various amounts of electrical current spread. I want to illustrate the differences between the traditional type and the cochlear implant simulation here for you…

Slide 17

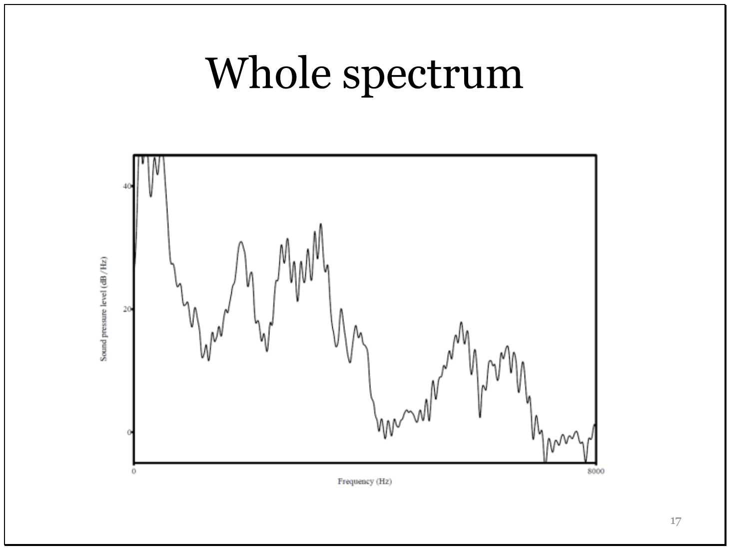

Typically you take the frequency spectrum,

Slide 18

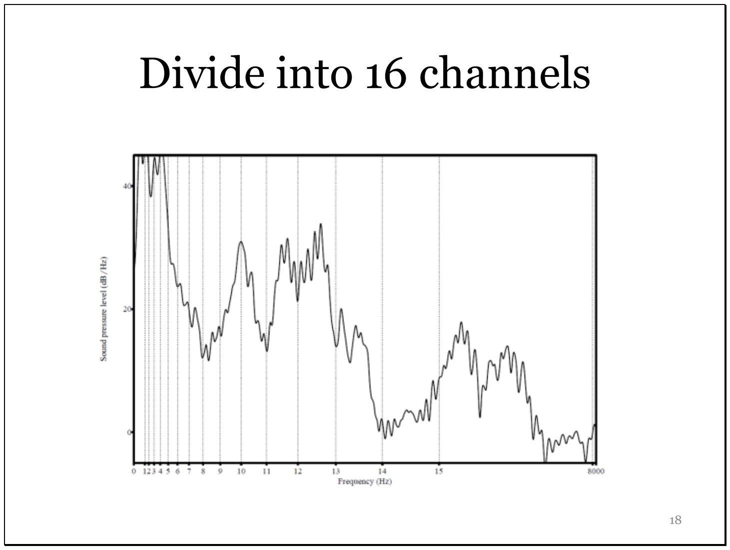

You divide it into a number of channels, say 16 channels

Slide 19

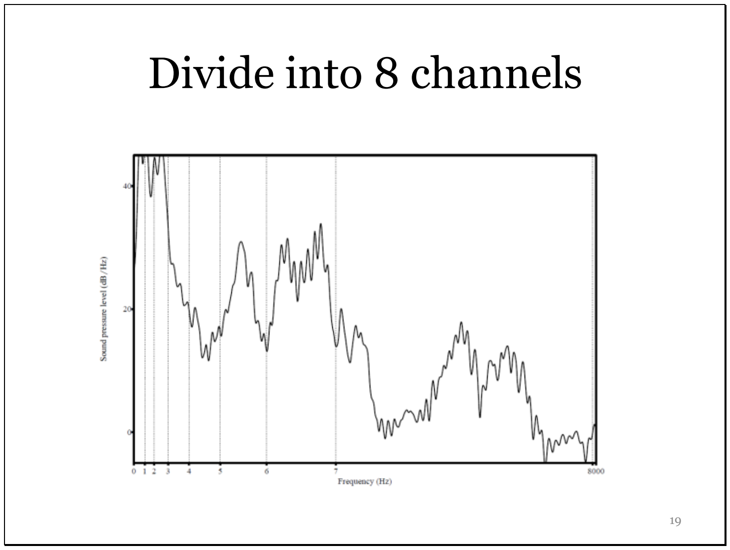

Or 8 channels.

Slide 20

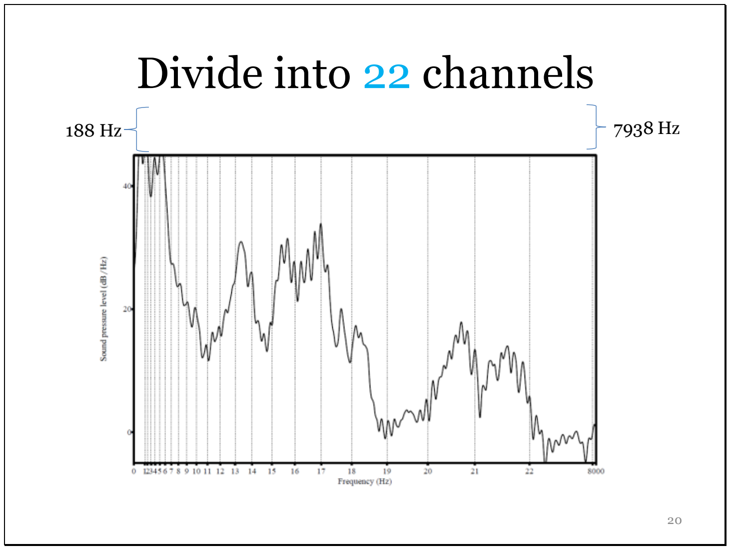

So what we did was we matched the channel frequency allocation to the cochlear device
Including the correct frequency range used in the Nucelus processors.
And if you think about the speech processing strategy that these listeners use, they don’t actually get all 22 at once,

Slide 21

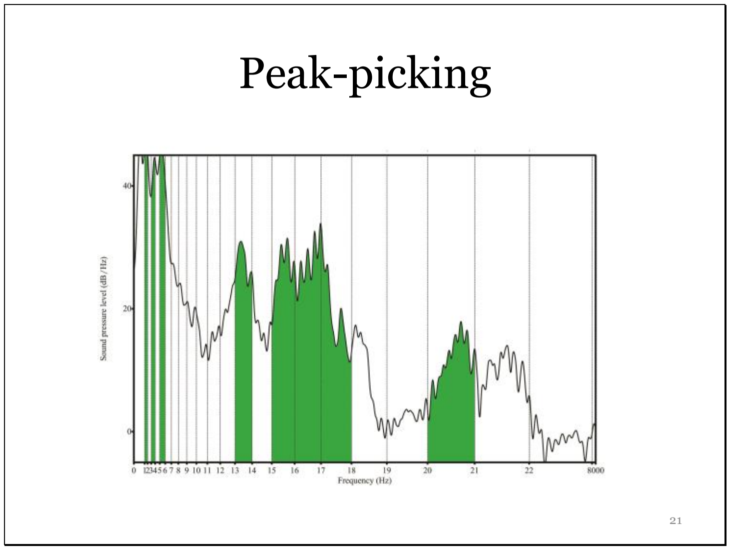

Out of those 22 channels, 8 get picked for activation at any one time.
But they aren’t so narrow and specific like this, or else CI listeners would all perform at excellent levels, which they do not.

Slide 22

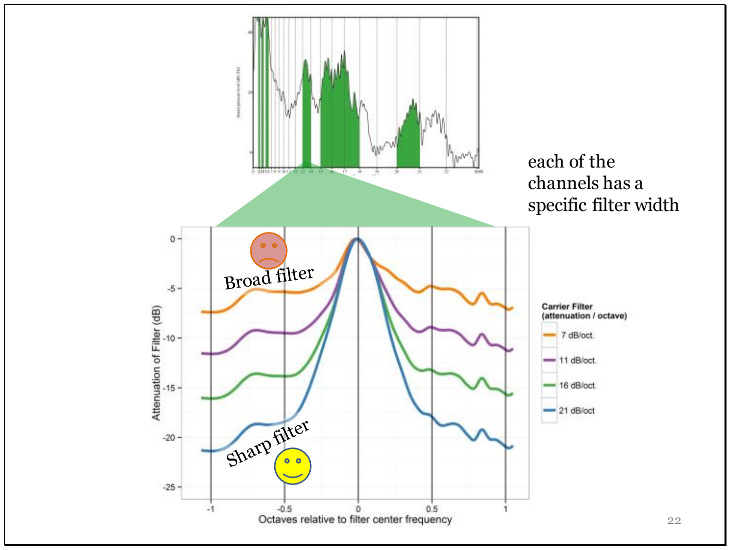

Instead, each of the channels has a filter width that can be more or less specific. Think of this as the amount of current spread from each electrode in the implant.
You can have a nice sharp filter for fine grained detail,
Or you can have a broad filter, smearing all the information and degrading the spectrum.
It's thought that people with cochlear implant have something more like that broad filter, or even worse, and *that's* why they perform poorly.

So, another kind of spectral degradation, this time more reflective of what we see in people with cochlear implants.

If we find the same kind of effect that we saw before, we have some new insight as to the impact of such spectral degradation.

Around 3000 Hz:
½ octave is 2.4 mm.
1.47 dB/mm
2.31
3.36
4.4 dB/mm

Slide 23

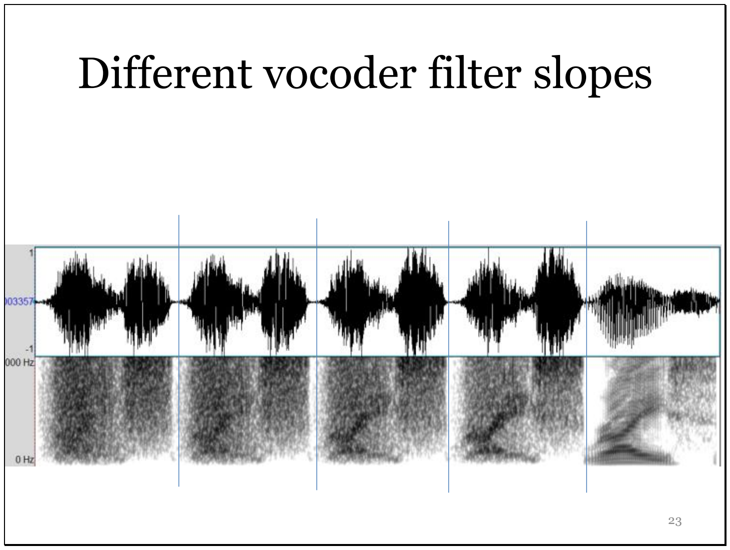

Here's a spectrographic representation of the different filter widths. Beginning at the left, we have the widest filters, with poorest resolution. At the right, we have the regular unprocessed speech. Each panel has the word "rice"

Slide 24

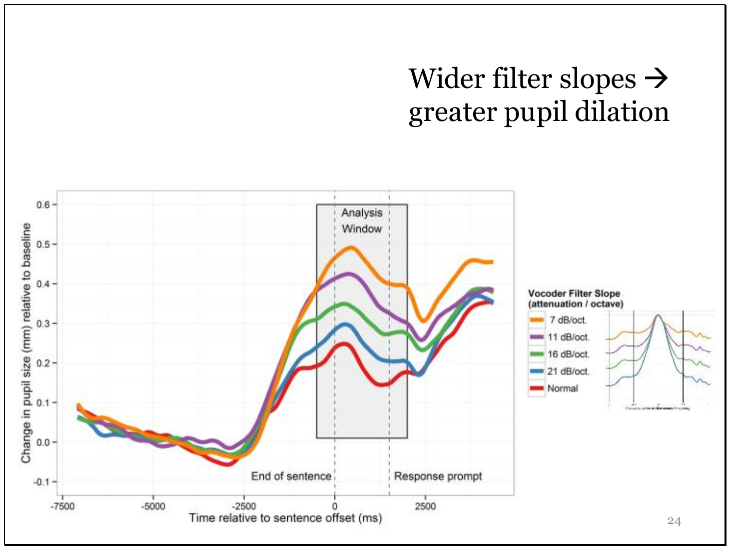

So when we simulate cochlear implant processing in this way, we get another nice systematic effect – similar to the one we saw before, except now we're being more faithful to what a cochlear implant actually does, and perhaps explaining something relevant to what they experience.

Slide 25

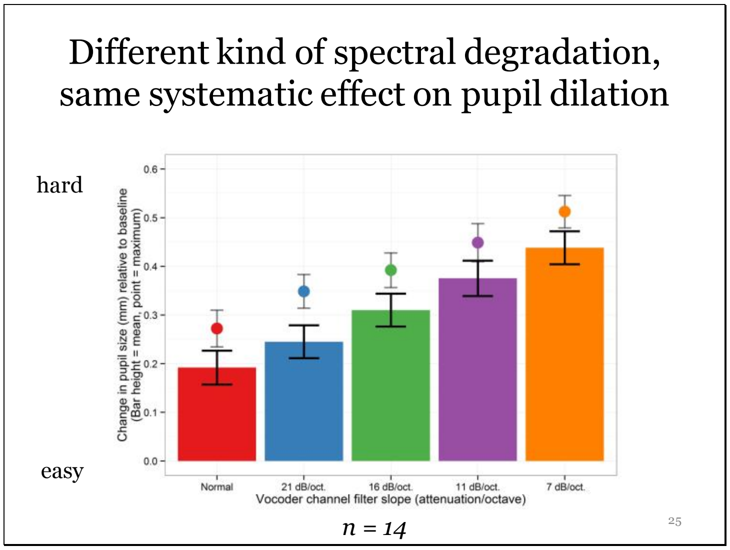

If you prefer seeing it in bar graph form, we can view it this way as well.

Slide 26

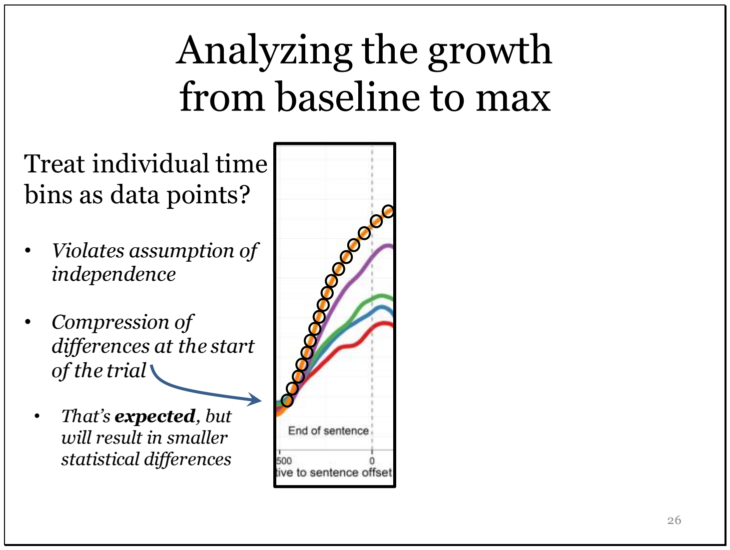

I'll remind you, we have this very fine granular measurement that is sampled over time, and we're comparing that measurement for our different conditions.
I'm going to walk you through how we've arrived at using growth curve analysis as our preferred method.
What we're doing is modeling the growth from baseline to maximum, which is this window here [].
We can treat each individual sample as a data point,
But this violates the assumption of independent samples that permits us to use good parametric statistics.
You also see that at the start of the time window, everything is compressed, right here.
This is expected because baseline is always at zero, but this shrinks the differences that we can clearly see are there, and thus weakens our analysis.

Slide 27

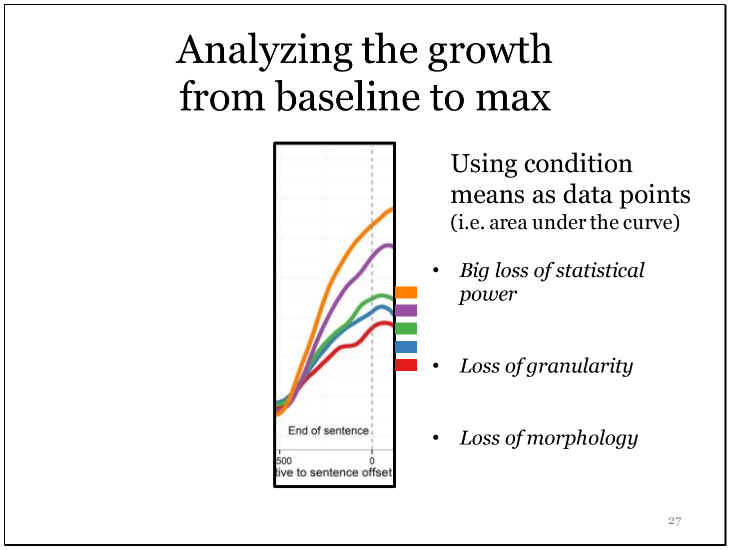

We can take each condition mean as a single data point, where you have compression at the start of the trial, but you at least don’t violate independence of data.
The problem is, you took ALL these time-sampled measurements, and you just got rid of all that granularity by dumping them into a single bin. You lose the morphology, and if it looked more complex than our simple curves, you’re really running the risk of losing true differences.

Slide 28

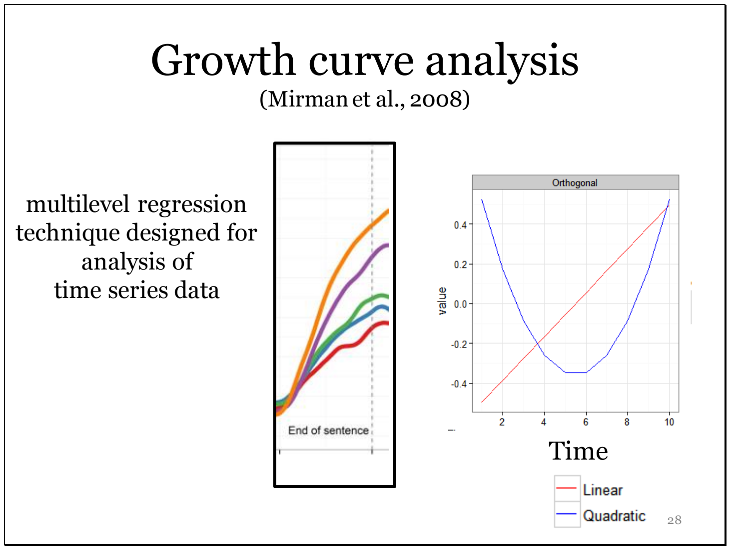

So our approach is to use growth curve analysis, which was described in nice detail by Dan Mirman a few years ago. This is designed for the analysis of time-series data. What you do is model the effect OF time on these conditional variables, and time can either move linearly or quadratically, which is simply to think of it as acceleration or deceleration.

The problem is, these two time curves are related to each other, so they steal the variance away fro each other and weaken your analysis.

So, a better solution is to use orthogonal time polynomials, which can go to quadratic, cubic, quartic, whatever polynomial best fits your data. In this case, linear and quadratic give a good fit, so we'll use that.

Slide 29

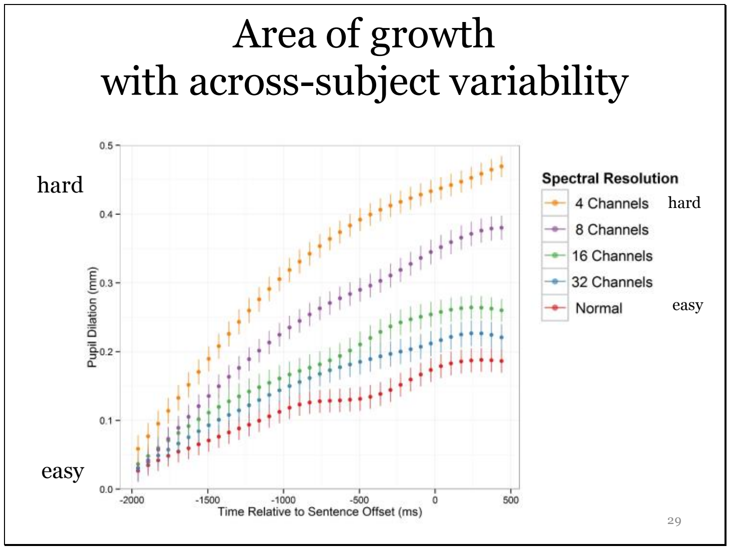

So let's take a closer look at this data. We see individual sampled data points, and the cross-subject variability, which was pretty tight.

Slide 30

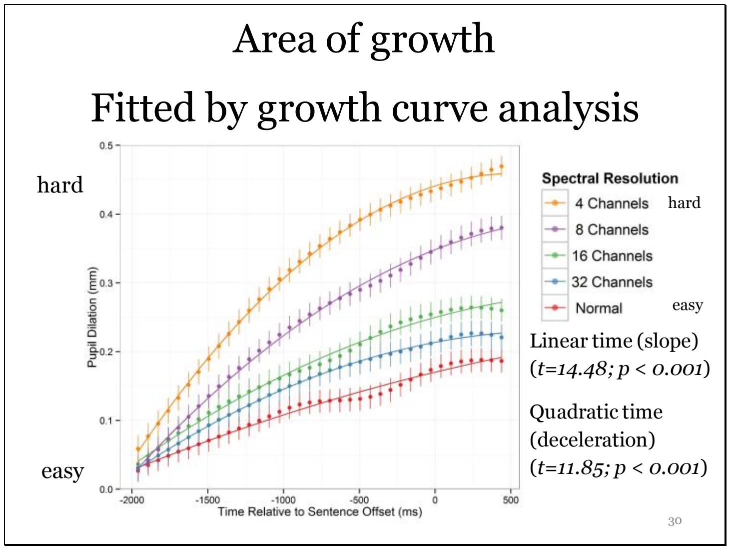

And now this is how the growth curve model fit the data – very well as you can see.
Now, as far as the statistics, we found a significant linear relationship between all the conditions here in terms of rate of change, or linear time.
There was also a significant effect across conditions in terms of quadratic time, or deceleration.
p-values were derived in a mixed model by assuming convergence of the t and z distributions.
In that statement, what I'm saying is that there is an ordered linear relationship between these conditions and the pupil dilation measured over time.
For those of you concerned with exactly which curves are different from the others – they all are different from each other.

In terms of acceleration all different except 32 & 16 and 16 & 8

Slide 31

Potential

- Individualized analysis for CI speech processing strategies that provide better spectral resolution

Test case

- Bilateral CI listeners using “interleaved” channel maps
  - (every other electrode is disabled, to cut down on interaction of adjacent electrodes)

Slide 32

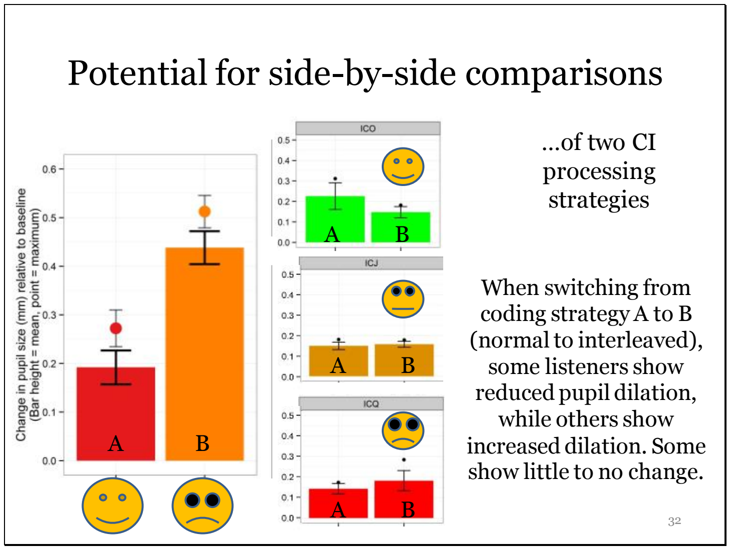

We have five different conditions here, but let's imagine for a moment that we want to simplify and just compare A versus B.
Put these two together, and now the important question is, which is easier?
And that is just what we as when we set up two different processing strategies or maps for a patient with a cochlear implant.
That's exactly what we did, and although I don't have time today to describe the details of the strategy, I can say that when testing it, a useful measurement has been pupil dilation. We see this listener here showed improvement that we can see with lower pupil dilation during listening…
And this listener showed decline he had to try harder when listening with our new map.
And, as is the case with almost all CI data, we have a third of the patients who go one direction, a third who go the other, and a third who show no effect.

Slide 33


Wide range of benefit / decline
ICO
ICG
ICD
IAJ
ICP
ICJ
ICF
IBK
ICM
ICQ
Normal
Interleaved
Change in pupil size (mm) relative to baseline
(Bar height = mean, point = maximum)
Map Configuration
Ratio
(Interleaved : Normal)
0.7
0.8
0.9
1.0
1.1
1.2
benefit
decline

Slide 34

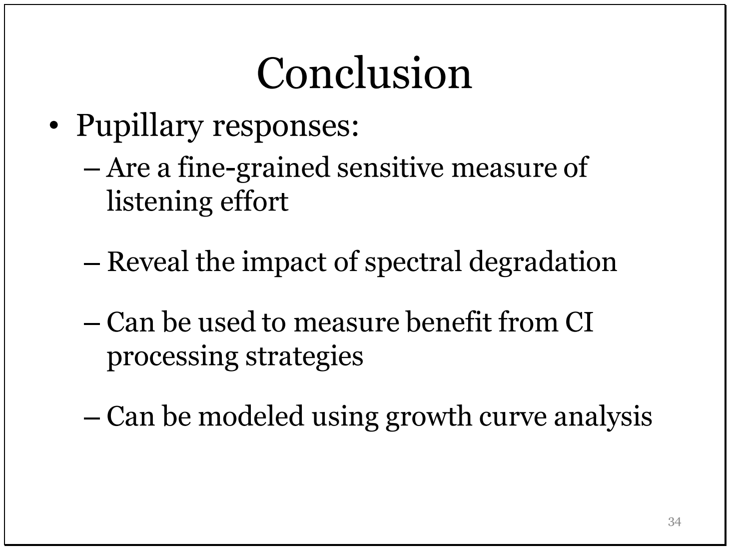

So to touch on the main issues we've seen today,
We've seen that pupil dilation is a nice fine-grained and sensitive measure of listening effort.
We see that spectral degradation has significant consequences in terms of listening effort, whether you do this with a traditional vocoder, or using a different vocoder that simulates real-life cochlear implant processing.
I gave you a little sneak peak at how this approach can be used to evaluate different processing strategies for patients with cochlear implants,
And finally, I've shown you how we can model the pupillary response using growth curve analysis, which is a robust and appropriate method for finding differences in time-series data.

Slide 35

I've worked on this project with Jan Edwards and Ruth Litovsky, whose grants have supported this work, which was also aided by a core grant where we work, at the Waisman Center.

Tristan helped write the scripts to handle the massive amounts of data generated by this project.

Rob helped write a real-time calibration procedure to control visual field luminance.

Kayla helped collect and score intelligibility data

Erica and Franzo both helped to implement growth curve analysis for these data.