Acquiring language:

A story about research

Micha Elsner, Department of Linguistics



Back in the early 80s...

- Baby Micha born (Jerusalem, Israel)
- Immediately starts acquiring language
- What was I really learning?



Before the 80s

- International Phonetic Society founded 1886
 - At first, English, French and German...
 - But later, sounds of different languages worldwide
- Electronic signal processing: 40s and 50s
 - Allowed detailed study of acoustics of speech
- Not much is known about early infancy...
 - Babies don't speak
 - Nor do they answer lab questionnaires

By the 80s, this is starting to change



Infants learn phonetics very early!

- Janet Werker and colleagues do an experiment...
- Infants listen to English/Hindi/Salish sounds



By 8 months, I'd learned some sound categories



- Including the 12(ish) vowels of English
- Probably also the 5 vowels of modern Hebrew
 - Which my parents often spoke until they left Israel a year later
- A few months later, I started to talk myself...

But Werker's result left researchers puzzled

- Phonetic learning begins *very* early
 - Before most of social cognition
 - Before infants can make the sounds themselves
 - Before knowledge of words and meanings
 - (In 1980, researchers think pre-verbal infants know few words)
- So how are they learning?
- By the mid-90s, researchers had come up with an idea...

Distributional learning

- Pay attention to *rare* vs *common* patterns
 - An idea drawing on Artificial Intelligence...
 - $\circ~$ And before that, from WWII codebreaking
- In 1996, Jenny Saffran showed infants can learn words from just two minutes of monotone audio!



<u>Stimulus</u> *ki-bu-*go-pi-*ki-bu*-la-ti-*ki-bu*...

So I went off to college...

- Majored in Computer Science
- "What's Linguistics? Will it fill my social science requirement?"

I'm going to develop artificial intelligence!





What's on their website





In my class on language acquisition

- Read a paper by Jessica Maye with Janet Werker and LouAnn Gerken, published 2002
 - Test distributional idea on *sounds* instead of words
- I didn't realize it at the time...
 - But this was cutting-edge research!

Linguistics is pretty interesting.

Maybe I can work on talking robots!



Maye teaches infants minilanguages





After a few minutes...

• Use the Werker setup to test perception



• Infants in group 1 detect the change better!

I passed the class, then didn't think about acquisition for a while

Instead, I got a job as an RA...



Eventually, they let me hack the parser a bit...

• We wrote a 4-page workshop paper....

Micha Elsner; Mary Swift; James Allen; Daniel Gildea Online Statistics for a Unification-Based Dialogue Parser

• And I started thinking about grad school...



Getting into a Ph.D program

- You are applying for a job as a researcher
- Make the case:
 - You know what research is actually like
 - You are independent and dedicated enough to do it
 - You have some interesting ideas to work on
 - Your interests are compatible with an advisor's
 - And with their grant funding!

So, your statement explains:

- Any research experience you have
 - Did you contribute your own ideas?
 - If not, why are you sure you'd be a good researcher?
- What you want to do next
 - And who you want to work with (mention names!)
- Anything that went wrong...
 - If you have a bad grade in a key subject, explain!
 - Is there evidence that you're better now?

Meanwhile, computer modeling steps in

- Test the limits of Maye's claim
 - Build a prototype distributional learner...
 - Show it works in her experiment
 - But can it learn real categories?
- de Boer and Kuhl (2003): yes it can!
 - Child-directed speech works better
 - \circ Only tried it for /a/, /i/ and /u/ :(



de Boer and Kuhl's learner: data

Vowels characterized by *formants* (resonances of the vocal tract)

• Since 1950s



Vowel data in two dimensions



Starting with an uninformed guess...



Sounds are probably members of the nearest category



Temporary confusion may arise



Continuing to shift the categories to fit the points fixes this



But I wasn't working on that...

I got really excited about *coherence* (relationships between utterances that make a discourse make sense)

And ended up studying internet chat rooms...

Who's talking to whom?



Brown University Computer Science



Eugene Charniak: my advisor

5.5 years in grad school

• Research starts immediately



- Also two-ish years of coursework
- But good grades won't save you from poor research
- When not doing your own research
 - Go to lab meetings and hear about other projects
 - Read papers and learn new techniques
- Many grad students also teach courses
 - But I was just a TA

At our weekly reading group...

Sharon Goldwater studies infant word learning:

- Built a Saffran-like model which learns 80% of words in written transcript
- No acoustics, though

Naomi Feldman studies sound categories:

- Working on Kuhl-like model for vowels
- Using fancy cutting-edge statistics
- But running into problems...





Why Kuhl's model doesn't work

"Our simulations suggest that this lower degree of overlap between categories may have been critical to the models' success."



Feldman's new idea

- Not just distribution of vowels overall
- Also ideas about lexical items
 - Infant hears "cat" but never "cet"
 - "let's" but not "lat's"
- By mid-2000s, clear that babies know some words by 6-8 months

Adding word learning helps



So, Maye is (a bit) wrong... distributional learning on its own isn't enough

Grad school: hard on mental health

(If you're having issues with depression or anxiety, your institution can probably help.)

- What you're doing often doesn't work
- It's not clear how to fix it
- You meet a lot of people smarter than you
- You set your own goals and schedule
- And just when you get good at it, they make you leave...





Non-academic options with a Ph.D.

Disclaimer: these jobs mostly for people who know code and stats

Industry: Google, Microsoft...

- Pros: More money for equipment and staff
- Cons: Less self-directed; more product development

Startups: Prismatic, Mixpanel...

- Pros: Live in San Francisco; work with small, brilliant team
- Cons: No job security; riches or ruins

Government: NIST, DARPA...

- Pros: Good pay and benefits
- Cons: Rarely doing the coolest research (except spies!)

Some fields also have clinical jobs (like Speech Therapy)

I wanted to stay in academics, so I got a postdoc

- Short-term mercenary researcher
 - Hired with grant money
 - Usually 1-3 years
- Career development:
 - Meet new contacts
 - Publish new papers

I just got a grant! You should apply for the job...

It's good to have contacts.





I'm already excited about acquisition

- By 2011, we believe:
 - Infants learn words and sounds very quickly
 - Early learning works by counting
 Bare vs common patterns
 - Rare vs common patterns
 - Learning words helps infants learn sounds
- But natural speech is full of variation
 - Sometimes "and", other times "en"
 - How can infants cope?

Started work with transcribed data

(Ok, some caveats about this data. We can discuss.) *y uw* || **w** aa n || *t uw* || *s iy* || *dh iy* || *b uh k* || "You want to see the book?" *I* uh k || dh eh r s || ah || b oy || **w** ih || ah s || hh ae t || "Look! There's a boy with his hat." eh n || ah || d ao g iy || "And a doggie!"

While debugging my model code, I stared at this file for hours every day...



Words and sounds



With variation, fewer bogus "words"

Words containing "you" from our model:

you (805 times), doyou (240 times), youwan (88 times), yih (58 times), areyou (54 times), youdo (47 times)

Words containing "you"; no phonetic variation:

you (498 times), yih (280 times), ya (165 times), yee (119 times),

doyou (106 times), doyee (44 times), canyou (39 times), canyee (29 times)

Our model learns a compact early lexicon

• More similar to real infants in the lab

Being a postdoc is awesome

Except:

- Only lasts 2-3 years
- First year spent moving
- Last year spent looking for a new job
- Hard on young families



Academic job options

"Senior researcher" positions

- Just research, no teaching
- Some are great, others are glorified postdocs with no job security
- "Adjunct"/"lecturer" positions
 - Teaching 24/7, often with no benefits
 - Usually terrible, and hard to escape, too!

Tenure-track positions ("professorships")

- A mix of research and teaching
- Good job security and pay
- Most academics want one, so they're hard to get

How faculty get hired

- Department begs for money to hire someone
- Job ad posted on professional websites
- You and 100-200 other people apply
 - Mostly new Ph.D.s, postdocs or junior faculty
- List is whittled down to 3-6
 - Based on recommendations, publication record, statements, how well candidates "fit" with department
- Finalists give a talk survive a 2-day

Tenure track: 6 year trial period

Priorities:

Research! Publish and get grants

Teaching: Don't mess this up

Service: Administrative work, committees and outreach (like this talk) **Not being a jerk:** Your colleagues vote on whether to give you tenure, so try to be nice

Applying for a grant

"Cognitive models of the acquisition of vowels in context" Micha Elsner (OSU) and Naomi Feldman (UMD)

- Planned work:
 - Model variation in acoustics, not transcripts
 - Test models on real child-directed speech in multiple languages
 - Extend models to handle more kinds of variation

We wrote up our plans... And sent them to the National Science Foundation

- Computational Cognitive Science panel
- Asked for 3 years of funding
 - Two Ph.D. student salaries
 - A one-year postdoc to help with Japanese
 - Travel to conferences

Once the grant is awarded...

Have to recruit students to help with the research...
 Currently working with Stephanie Antetoma



Currently working with Stephanie Antetomaso and Martha Austen... neither of whom are rhesus monkeys





My current work...

Adding acoustics to previous word learner:

Then:

y uw || **w** aa n || *t uw* || *s iy* || **dh iy** || *b uh k* ||

Now:

y <380.5 1251.6> || w <811.8 1431.9> n || t <532.9 1094.1> || s <468.2 2703.2> || dh <595.2 973.8> || b <545.3 1330.0> k ||

• Will have to deal with more realistic kinds of phonetic variation which previous models ignore

Testing on real data

These vowels are from carefully controlled speech:

"had, hod, who'd..."

In general speech, things are messier





Better tools for getting real data

You've already seen that transcribing is hard

• Which means it's expensive!

So I'm also working on better tools for automatic transcription

Current open questions

- How can we model phonetic variability to cope with real, messy data?
 - Do we have to model specific kinds of sound changes, like reduction (*and* to *an*) separately?
 - Are learners predisposed to deal with some changes better than others?
- How does learning about variability work for bilingual or bidialectal children?

Plan your career with a little cynicism

THE EVOLUTION OF INTELLECTUAL FREEDOM



WWW. PHDCOMICS. COM

But somehow, we do make progress





Questions / Discussion