UC Berkeley
Berkeley Scientific Journal

Title
Tackling Language Learning: One Model that Rules (Dr. Steven Piantadosi)

Permalink
https://escholarship.org/uc/item/02w063b4

Journal
Berkeley Scientific Journal, 27(1)

ISSN
1097-0967

Authors
Guan, Grace
Niyogi, Anjuli
Palayur, Ann
et al.

Publication Date
2022

DOI
10.5070/BS327161282

Copyright Information
Copyright 2022 by the author(s). All rights reserved unless otherwise indicated. Contact the author(s) for any necessary permissions. Learn more at https://escholarship.org/terms

Undergraduate

eScholarship.org Powered by the California Digital Library
University of California

Tackling Language Learning: One Model that Rules

Interview with Dr. Steven Piantadosi

By Grace Guan, Anjuli Niyogi, Ann Palayur, Allisun Wiltshire

Dr. Steven Piantadosi is an assistant professor of the Department of Psychology at the Helen Wills Neuroscience Institute at Berkeley. He is also the principal investigator of the computation and language lab, which emphasizes its research on language acquisition and human conceptual systems. Some of Piantadosi's past research has focused on the ambiguity of language and the optimization of word length for communication. In this interview, we will be discussing one of his more recent papers, "One model for the learning of language." This paper describes a computational learning system that can acquire the key structures of a language, given enough positive data from the language itself.

INTERVIEWS

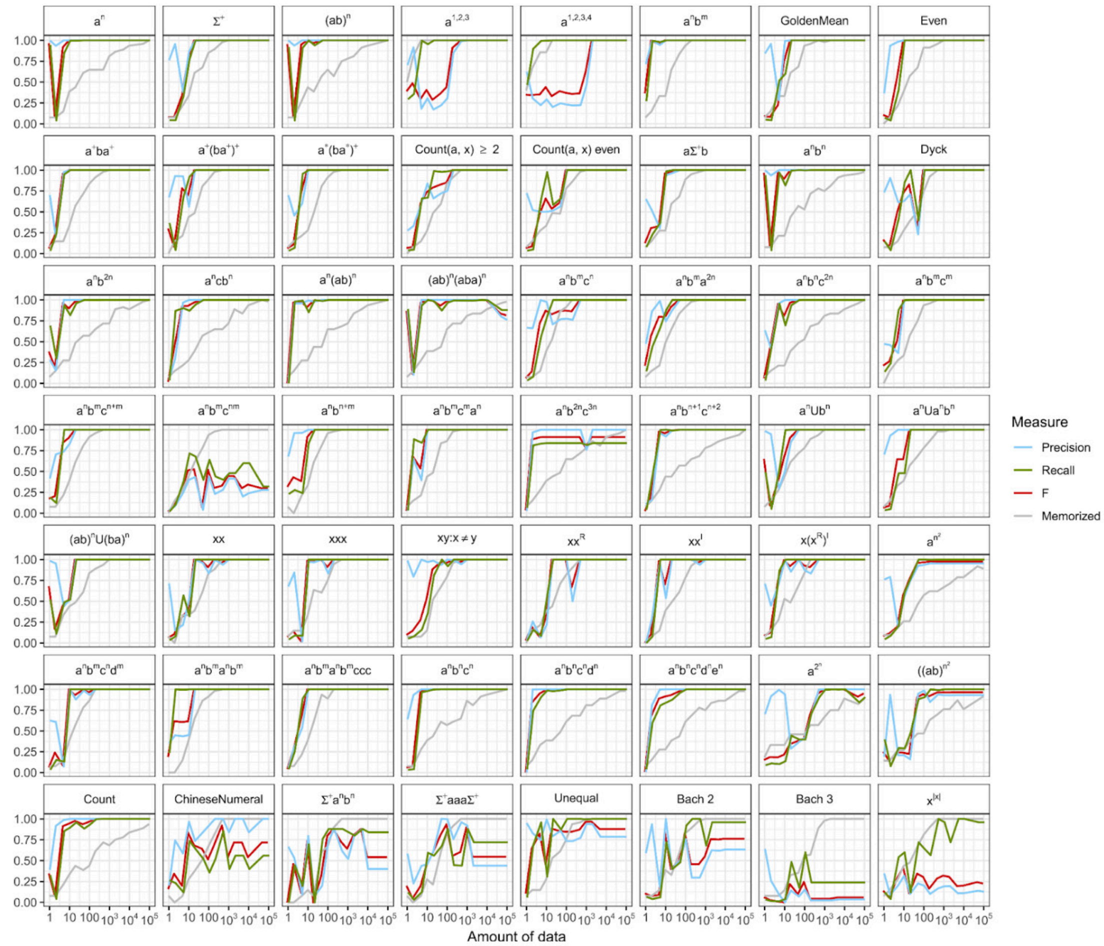

Figure 1: Fig 1 displays various possible language patterns made when forming words (i.e. a^n or ab^n, where a and b are arbitrary phonem). *For each pattern, the model received positive evidence of the data, As the amount of data fed into the program increases (x-axis), the accuracy of precision and recall generally increased (y-axis). For the memorization factor, where the system merely memorized the data inputted, the F score is performed much more slowly and poorly.*

BSJ: How did you become interested in cognitive science and language?

SP: When I was in high school, I read *The Language Instinct*, a book by Steven Pinker about the nature of language and where it comes from. It makes you think about the relationship between humans and other species as well as distinct human behaviors, such as our ability to think or learn in the world. I went to the University of North Carolina for undergrad, and they did not have much of a cognitive science program there, so I ended up studying linguistics and math, which ended up being a really useful foundation. When you put both of them together, you can get neat theories of cognitive or language behavior.

BSJ: In your paper, "One model for the learning of language," you discuss a learning model adapted from data of formal languages. How would you describe Gold's Theorem and its role in this study?

SP: Gold's theorem is this very old result. It assumes that when someone learns language, they use only what they hear to figure out what sentences are allowed and not allowed in the language. People started to analyze that kind of problem mathematically and ask, "Under what conditions would kids be able to figure out their parents' grammar?" Gold's result was an early finding in that spirit of mathematical analysis that shows, under certain mathematical assumptions, it is not possible for kids to pinpoint their parent's language from just observation, meaning that they couldn't figure out the rules of how language works.

Gold's theorem was really influential in linguistics theories about learnability. People took it to mean that language learners had to be highly constrained in the set of things they would acquire. In other words, most of the rules of language had to just be "built in" innately in order to solve this challenge for learning. They use that kind of mathematical analysis to try to deduce what human babies must know in order for language acquisition to succeed.

BSJ: How would you explain the learning model you created in its simplest terms?

SP: The learning model is kind of the opposite of Gold's Theorem. It comes from the point of view that there are lots of things humans learn in lots of different domains. There are simple procedures any human can learn, like tying our shoelaces or checking out at the grocery store. We can also learn complicated tasks such as how to take derivatives or build model rockets. If you think about Gold's Theorem of learning in the context of everything that people can learn to do, I think it becomes clear that the Gold approach is just not the right approach. Gold would say that you cannot really learn much of anything, much less the space of possibilities that humans seem able to. There are other mathematicians who have thought about this problem and have worked out learning theories that operate over the least restricted space possible: the space of computations. Think of every possible computer program you could write down; you could write a computer program that would tie shoelaces, another one that would check out at the grocery store, compute derivatives, etc. All of that human knowledge can be captured in the form of a program. A good way of making predictions about the world is to try to figure out how it works, and then use your theory to make predictions. The learning model is an implementation of ideas. It is one where people had previously thought about problems mostly in theory to determine how to build a really smart learner or a really smart AI system, but there were no implementations of that kind of learning. With this model, we were trying to cover a bunch of examples in linguistic research and show how a very general program learner is able to acquire them.

> "A good way of making predictions about the world is to try to figure out how it works, and then use your theory to make predictions."

BSJ: Could you explain how positive evidence is defined within the context of this study?

SP: Positive evidence refers to how individuals figure out the rules of a language by listening to the sentences of those around them. That is one of the key assumptions of Gold, but it is one that is probably not accurate at all. When kids are learning language, they are also producing language and learning how what they say is interpreted. If, for example, a kid intends for a ball to be placed on a stick but says the wrong thing—for example, "Put the stick on the ball"—and you place the stick on the ball, they will be able to see whether what they thought was going to happen actually happened.

Gold assumed positive evidence just because it makes it easier to analyze this problem of language learning mathematically. Everybody knew that it was not right, but we also wanted to assume it because we thought that learning could still work even under that assumption. We hoped that just looking at children listening to the sentences spoken by one's parents would say enough about language learning, despite ignoring all of the other kinds of input kids receive.

Science explained!!

1. Recursion: the process of applying the same function onto an input until a base case is reached
2. Building little data structures: can be done in programming to build, organize, store, and access data more easily
3. "if" statements: a conditional statement that tells a computer to only execute something if the statement is true

BSJ: Could you briefly explain what Fleet is? How was this library used to develop this model?

SP: Fleet is a programming library that allows you to specify a programming language, such as Python, Scheme, or C++. After you tell the core operations of the programming language, it takes your data and tries to figure out what program generated the data. In our case, we gave it a collection of very basic types of operations that people think are plausible for small computers. We do not give it everything in the programming language of Python. Instead, we give it the ability to do common techniques in computer science like (1) recursion and (2) building little data structures and (3) executing "if" statements and other similar operations. Then, we give it positive evidence from different string patterns that occur in language and ask it to find the program. It is the implementation which is solving this sort of abstract version of this learning problem.

BSJ: In your paper, you state that "Fleet implementation includes examples in other domains outside language,"[1] such as rule number learning and logical rule induction, showing that this approach of learning is generalizable across many domains. In what areas can Fleet be used beyond its current domains?

SP: Beyond its current domains, learning programs like Fleet can be used to model number learning and rule learning in general.

My lab has done a lot of studies on number learning in humans. There are many developmental experiments about what exactly kids know about a number of words, at what age they know them, how they transition through different levels of partial understanding of how numbers work, and what number words mean. Numbers are interesting because they are a case where you can see that what a two-year-old knows is very different from what a six-year-old knows.

We have used learning programs similar to Fleet to try to model number learning. In that setting, you first hear number words being used and then you have to figure out the concept of counting. This is a little procedure that takes a set and gives it a label based on how many things are in it. We have number learning models like that.

We have also developed models for general rule learning in humans. An example would be putting people into an experimental setting and having them try to discover a new rule that they did not know before. Sometimes, there are simple rules, such as word learning. For example, I could tell you that there is a new word called "daxi," and you would have to figure out what "daxi" means. Through trial and error, you may eventually discover that an object is "daxi" if there is another object on the screen that is larger than it. We have developed models for rule learning similar to this.

Systems like Fleet are nice because they show that there are very general learning models which govern many different types of learning; such as learning numbers, words, or other logical rules.

All of this is related to a bigger debate in linguistics about how specific our knowledge is to language. There are many kinds of rules of language, which leads to this debate about what is innate for humans and how much of the rules are innate versus learned.

BSJ: In your paper, you talk about visualizing the results of your study using precision and recall. Could you describe the dynamics between precision and recall?

SP: The dynamic between precision and recall demonstrates whether you are overgeneralizing or undergeneralizing. High precision and low recall values give a model that undergeneralizes, while high recall and low precision values give a model that overgeneralizes.

Humans usually generalize in some way because they can only keep a limited set of data in their brains. However, Fleet is able to learn the rules of the language because we can feed an infinite set of data to construct these models.

More generally, the way we are generating models using programs like Fleet is to take data and come up with a concise computational description of that data. That concise description will then apply in new settings where we have not used it before.

Oftentimes, it turns out that the most concise computational explanation of the data you see also generalizes to things you have not seen before. This is analogous to us humans: when we develop an intuition of a pattern, it allows us to notice it right away, even when encountering it in a new situation.

> Science explained!!
>
> ***Bach languages***: languages in which words contain letters with an equal number, "n", of occurrences in any order within the word (e.g. aacbbc or daadcc); they are generally context-sensitive[2]

BSJ: What languages was the model unable to learn and why?

SP: In Figure 1, we can see where the model does a crummy job. In particular, some of the plots at the bottom of the figure for the Bach languages are where the model does not do well. These are ones that are really difficult to express with the model's built-in programming operations. Generally, the languages that the model does not do a good job of are the ones that would take a really long program to describe. Since we have finite computational resources and a finite amount of time for which we can run this search, we ended up not finding good programs for those.

I think there is certainly something different in how people solve these kinds of problems compared to how the model does. We do not really know how people do it. But, even our crummy implementation here is able to learn most of the things that we would want to learn. Our main point with the model is just that the learning problem is solvable.

BSJ: What are the limitations of the model? How do you feel it could be improved?

Figure 2. The diagram below is a finite state machine, which describes an infinite set of strings. *Beginning on the state "0", there are an infinite number of walks one could take to form a valid string, including {tpst, vvs, vxv, …} given this graph theory. The diagram to the right is a specific finite state machine for English modal verbs (may, might, can, etc.) and similarly illustrates syntactically valid sentences that can be formed with these rules.*

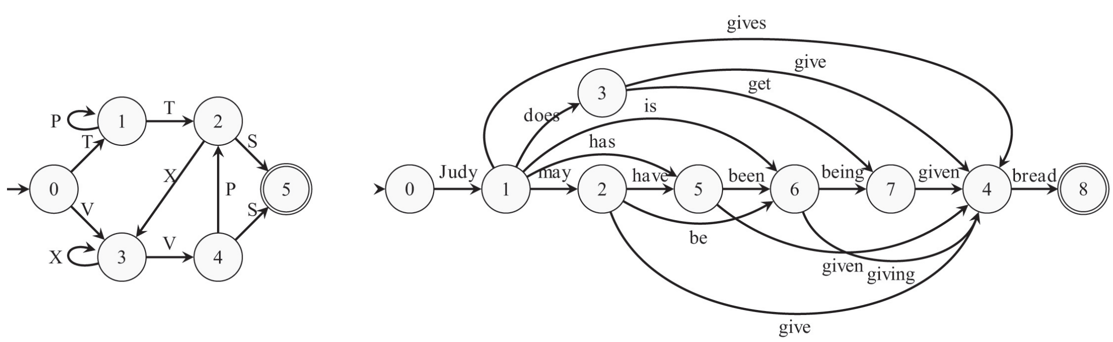

Morgan & Newport

S → AP BP (CP)
AP → a (d)
BP → CP f | e
CP → c (g)

Morgan, Meier, & Newport

S → AP BP (CP)
AP → o a (d)
BP → a CP f | u e
CP → i c (g).

SP: There are some languages that the model does not do a great job of learning, which may not be a real limitation because it may happen to be that people also do not learn those languages very well. Earlier, I was saying what makes the model good is that it shows high accuracy, meaning it can learn languages well. As a cognitive theory, however, what would make it "good" is if it learned languages in the same way that people did, meaning it makes the same kinds of mistakes and has the same strengths and limitations. However, we would need additional data to evaluate this speculation, such as connecting specific model predictions to empirical data. This is actually something we have done in the lab but have not published. It is a completely different way of thinking about what makes the model "good." Another limitation of the model is the search process, which is the model's ability to search through different programs to explain the data. It is similar to a random search; it is interesting that such an inefficient search process works well in learning these languages, which linguists have long argued are impossible. Humans likely have smarter and more efficient search techniques than the model uses.

BSJ: How does a model like this contribute to other fields in cognitive science?

SP: There is a big, underlying question in cognitive science, as well as neuroscience and education: "How does learning work?" We do not know the mechanistic underpinnings of real language learning, which is central to how we approach childhood education or knowing what kinds of educational interventions might be most useful. In the broader setting, I think of this work as really aiming to provide some kind of mechanistic understanding of learning. I work a lot with a collaborator, Jessica Cantlon who has a primate lab, and we are interested in what makes humans different from other species cognitively. It connects to human evolution, anthropology, and cross-cultural comparisons. I think of this model in that context, but I do not think any particular project is a very big step towards answering these broader questions yet.

BSJ: How does your research bridge evidence from different disciplines, such as philosophy, psychology, linguistics, and cognitive science?

> "Answering the question of what makes humans special requires you to draw on a bunch of different fields."

SP: This question of what makes humans special has been tackled by all of those fields. There are claims from linguistics about particular linguistic structures or processes being innate. I think this study is, in some sense, questioning those assumptions and claims. It is saying that maybe there is a more innate, fundamental attribute than those structures previously suggested, which is this ability to look at data and come up with some kind of computational account of it. There is the "child as a scientist" view, which has been written about a lot by Alison Gopnik, another professor here in the department, who bridges scientific inference, children's thinking and theory building, which goes in both directions. Answering the question of what makes humans special requires you to draw on a bunch of different fields. There are many techniques from computer science, for example, which are relevant here. Many underlying ideas about program learning are really from artificial intelligence, and people trying to figure out how to do induction from data. I think all of that is necessary for understanding how children learn.

BSJ: Are there any questions that this study has left you with? What do you hope to see in store for future research on language models?

SP: We have many questions that we work on in the lab, which are related to this study. Some of our research is looking for behaviors that differ between humans and primates, and behaviors that are shared across highly diverse human groups, including indigenous Amazonians. Some of our research is about connections between this kind of learning and other domains in cognition, like number, problem solving, or other parts of language. We are also interested in experimental work testing computational models: can we quantitatively evaluate the predictions that this kind of model makes? Resolving some of these basic mysteries of cognition will require an interdisciplinary outlook going forward.

REFERENCES

1. Yang, Y., & Piantadosi, S. T. (2022). One model for the learning of language. *Proceedings of the National Academy of Sciences*, 119(5). https://doi.org/10.1073/pnas.2021865119
2. Joshi, Aravind K, et al. "The Convergence of Mildly Context-Sensitive Grammar Formalisms." *Foundational Issues in Natural Language Processing*, Jan. 1990.
3. Headshot: [Photograph of Steven Piantadosi]. Received via email 12/1/22.


hello
SALUT
你好
¡HOLA!
bonjour
مرحبا