Are Neural Networks Good Linguistic Models? Then and Now

Sarah Brogden Payne
sarah.payne@stonybrook.edu
paynesa.github.io


Teen Academic Linguistics Conference
August 30th, 2025

The Million Dollar Question

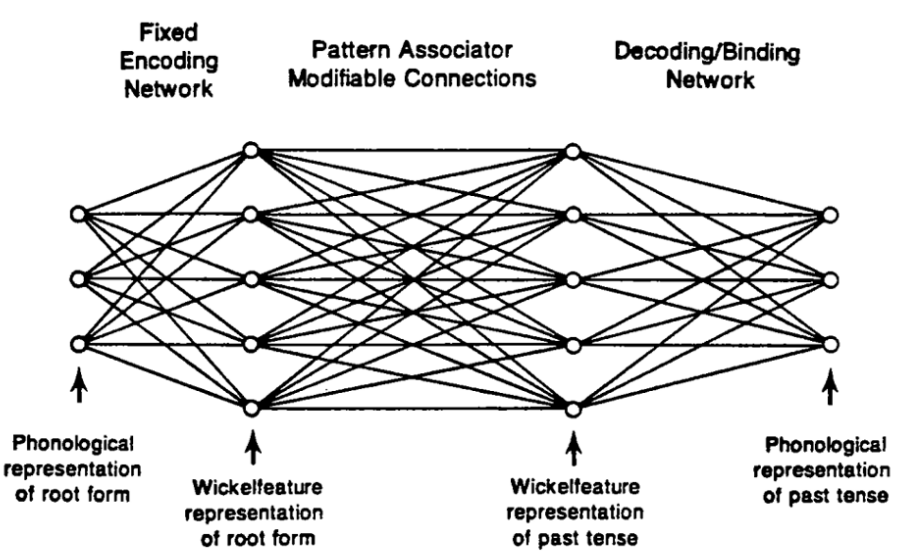

=

?

The Million Dollar Question

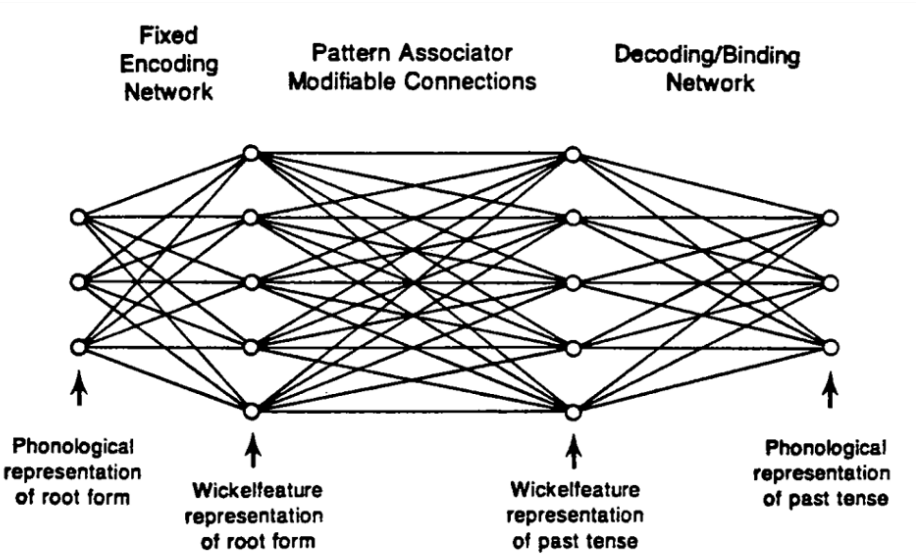

=

?

Why do we care?

The Million Dollar Question

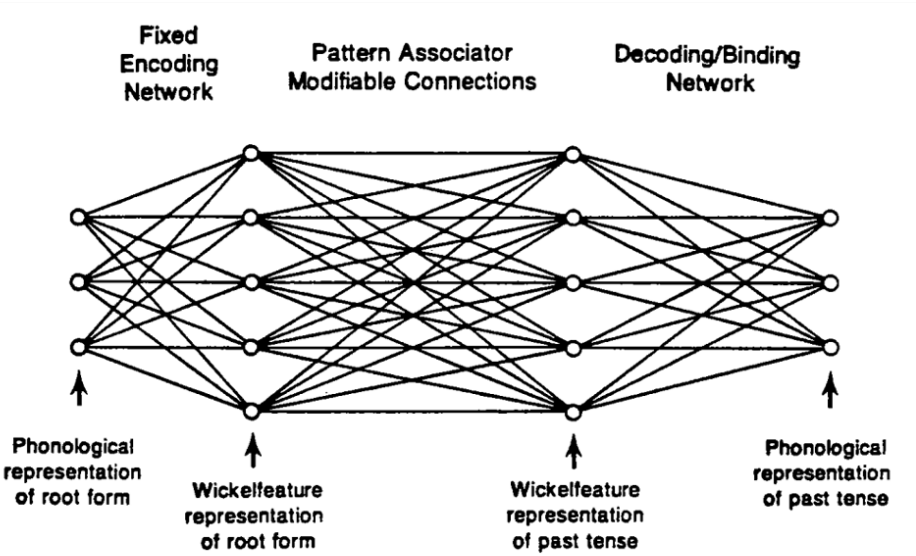

=

?

Why do we care?

Neural networks used by **connectionists**, cognitive scientists who believe:

The Million Dollar Question

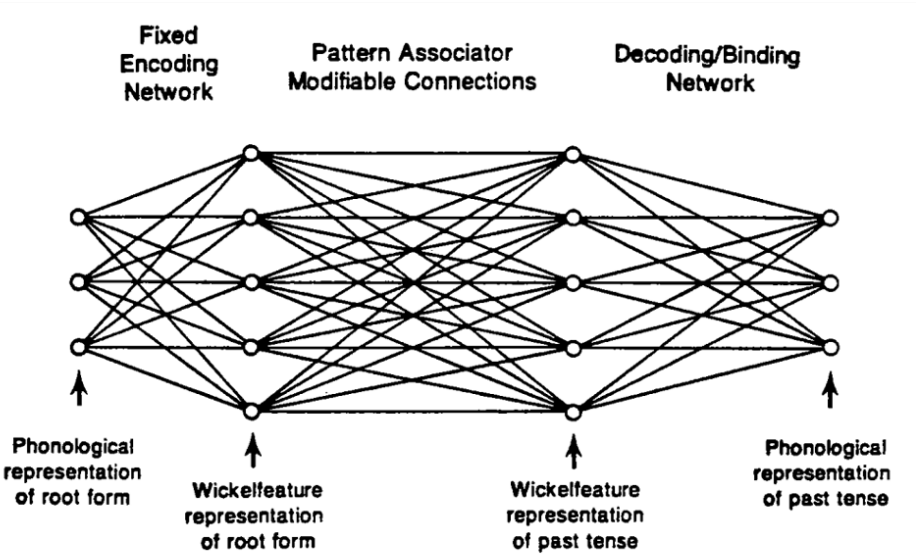

=

?

Why do we care?

Neural networks used by **connectionists**, cognitive scientists who believe:

- Representations are **subsymbolic**, distributed over networks of connected, simple units

The Million Dollar Question

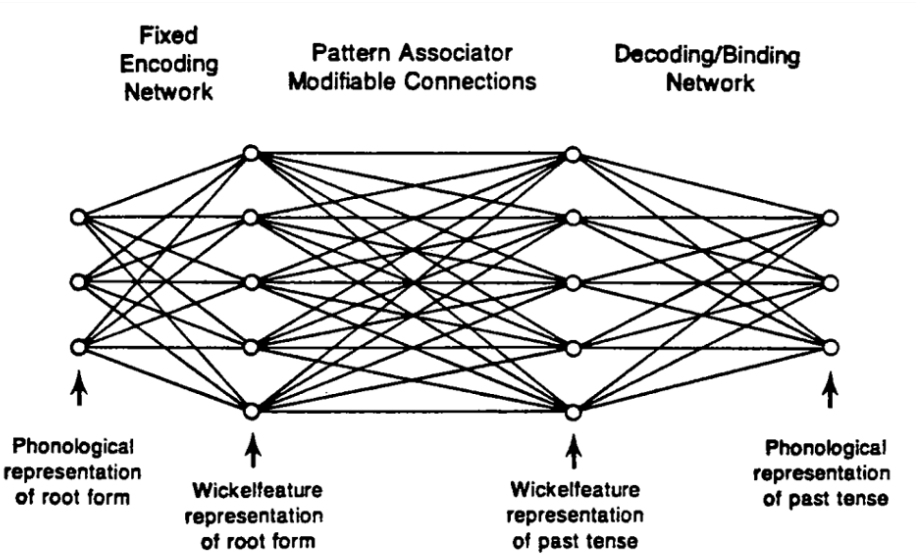

=

?

Why do we care?

Neural networks used by **connectionists**, cognitive scientists who believe:

- Representations are **subsymbolic**, distributed over networks of connected, simple units
- Everything is processed by the same **associative memory** mechanism

The Million Dollar Question

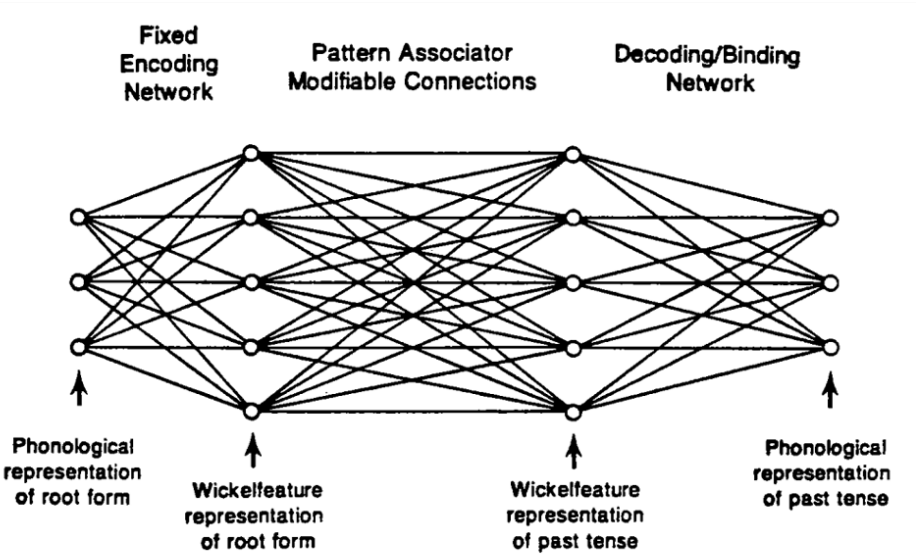

=

?

Why do we care?

Neural networks used by **connectionists**, cognitive scientists who believe:

- Representations are **subsymbolic**, distributed over networks of connected, simple units
- Everything is processed by the same **associative memory** mechanism
- Information is processed **in parallel**

The Million Dollar Question

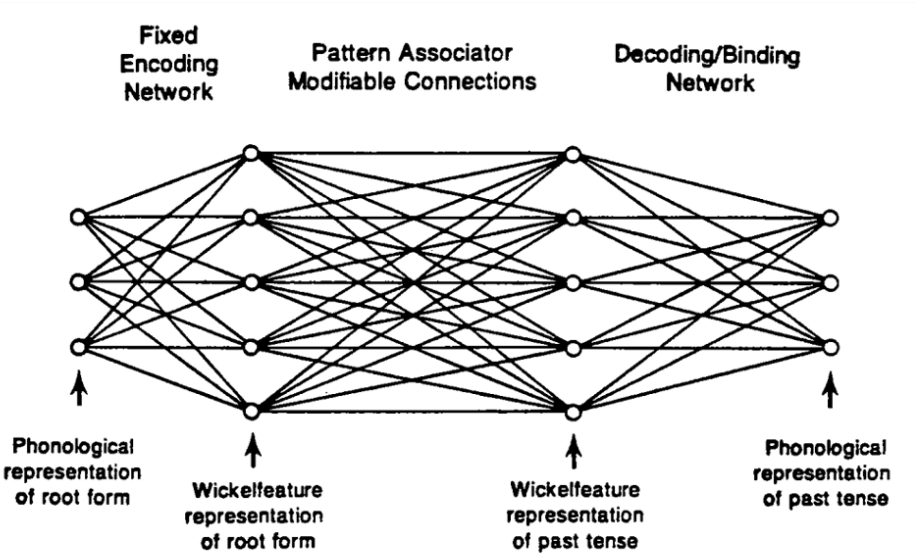

=

?

Why do we care?

Neural networks used by **connectionists**, cognitive scientists who believe:

- Representations are **subsymbolic**, distributed over networks of connected, simple units
- Everything is processed by the same **associative memory** mechanism
- Information is processed **in parallel**

Understanding NNs helps us understand **connectionist theory!**

The Million Dollar Question

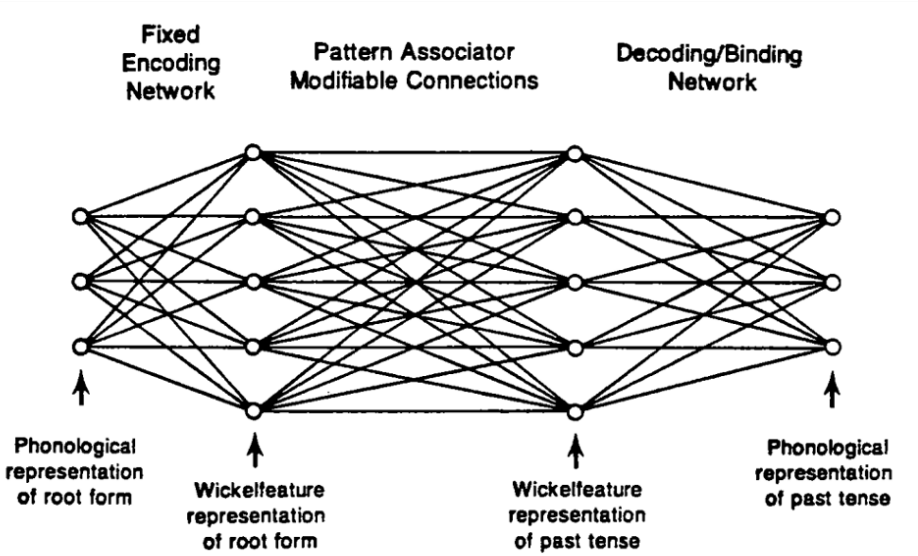

=

?

To answer this, we will need an understanding of:

The Million Dollar Question

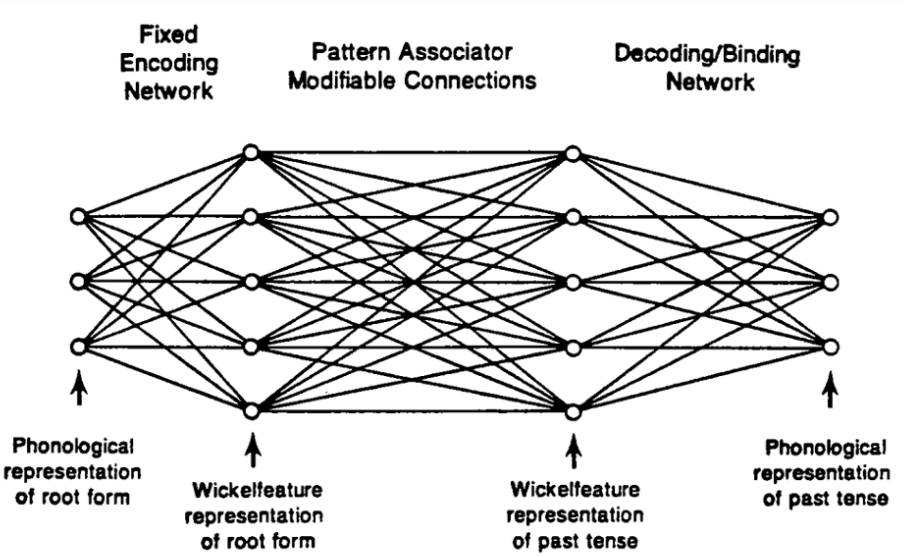

=

?

To answer this, we will need an understanding of:

- The **input** to the child learner

The Million Dollar Question

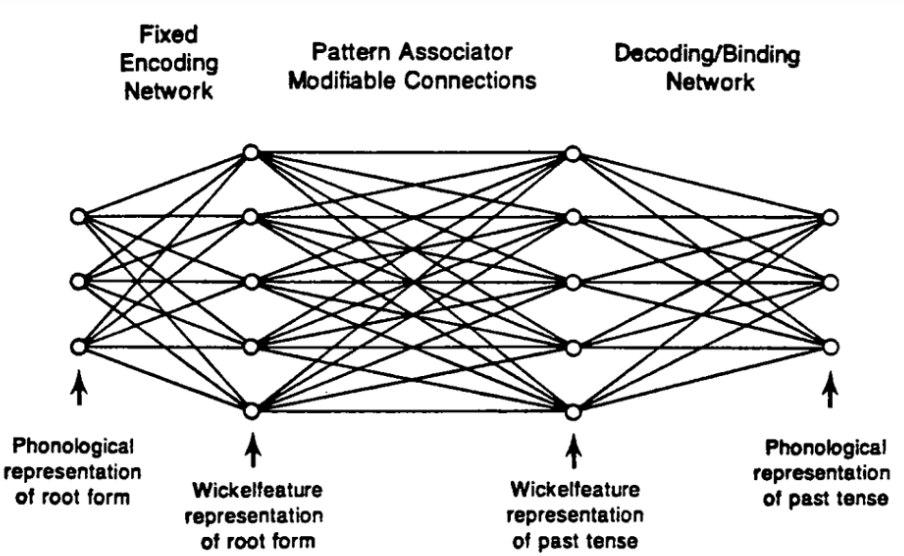

=

?

To answer this, we will need an understanding of:

- The **input** to the child learner
- The **errors** made by the child learner

The Million Dollar Question

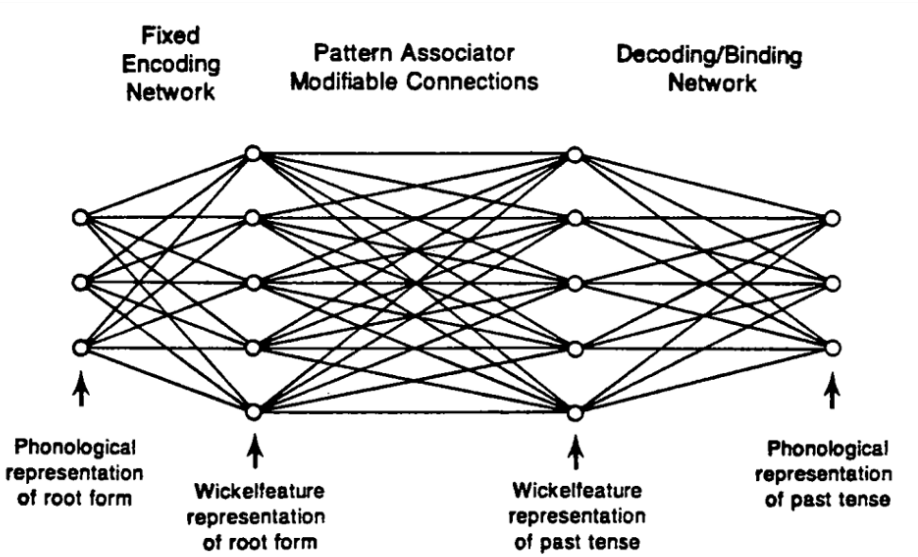

=

?

To answer this, we will need an understanding of:

- The **input** to the child learner
- The **errors** made by the child learner
- The **learning trajectory** of the child learner

The Million Dollar Question

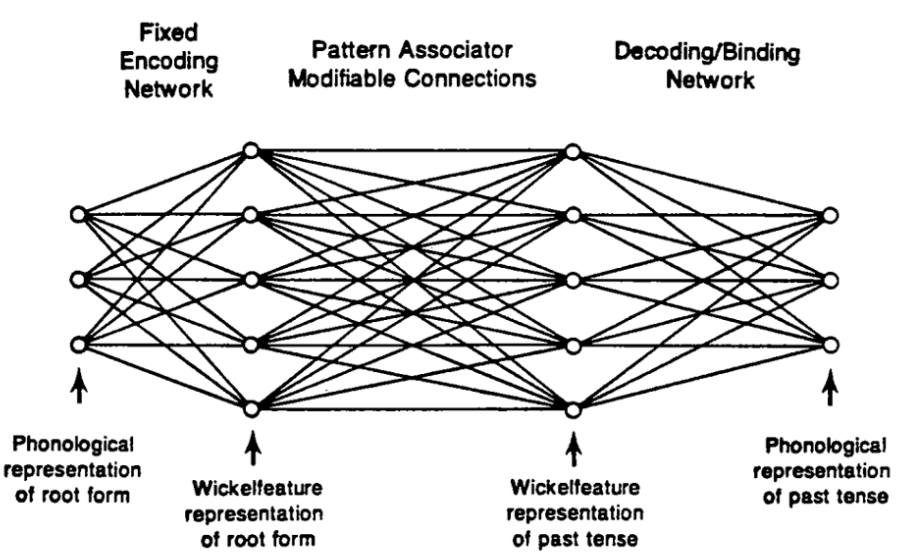

=

?

To answer this, we will need an understanding of:

- The **input** to the child learner
- The **errors** made by the child learner
- The **learning trajectory** of the child learner

How are these **the same** or **different** to those made by neural networks?

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Input

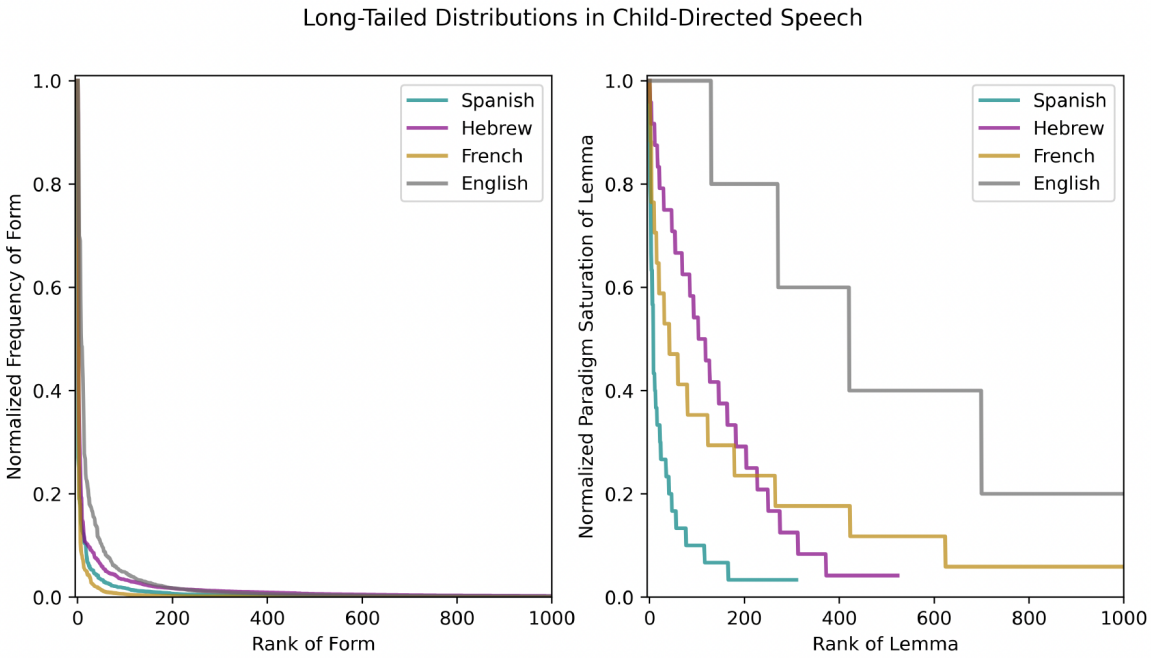

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Input

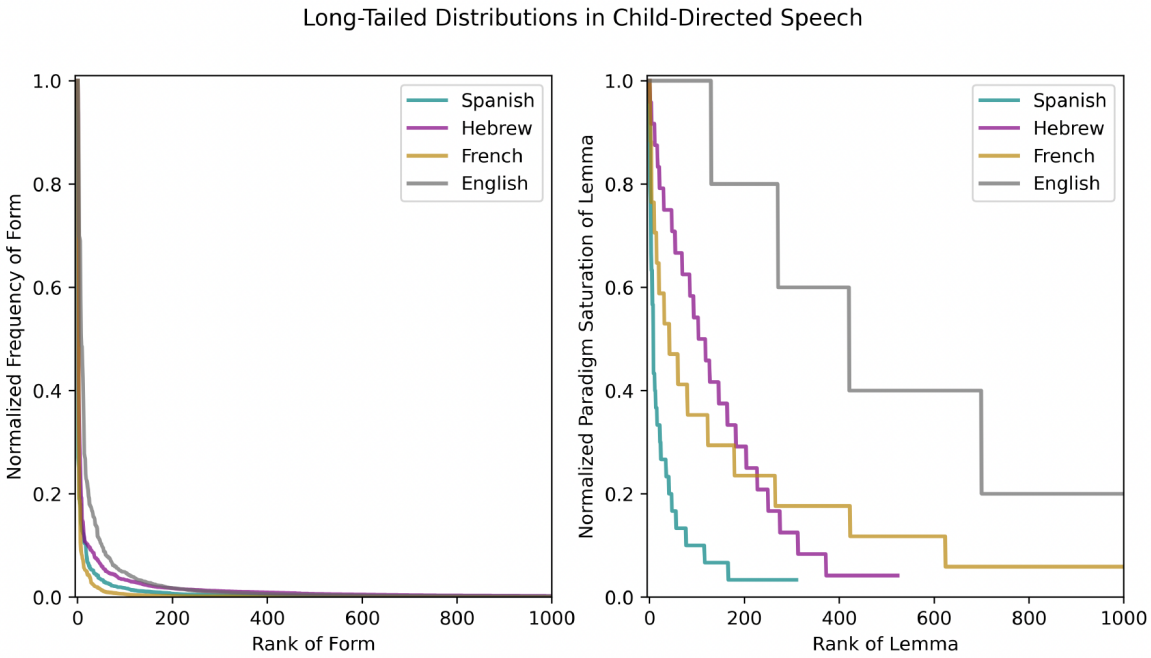

$\leq$ 1000 **words total**

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Input

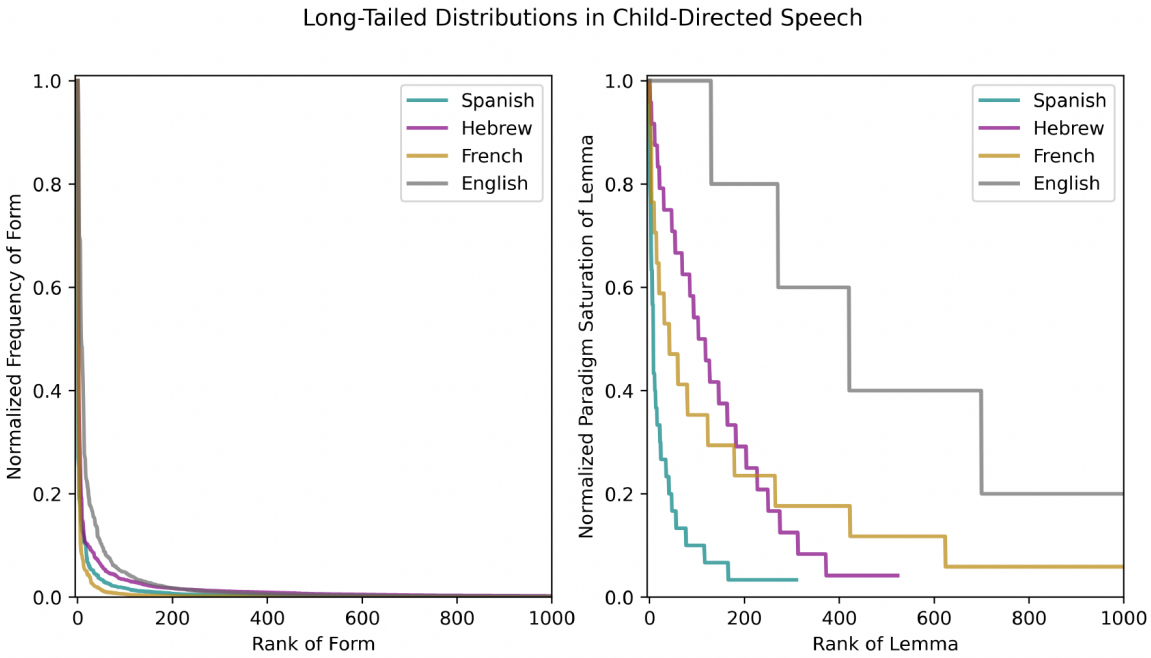

$\leq$ 1000 **words total**

Most words are **very infrequent**

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Input

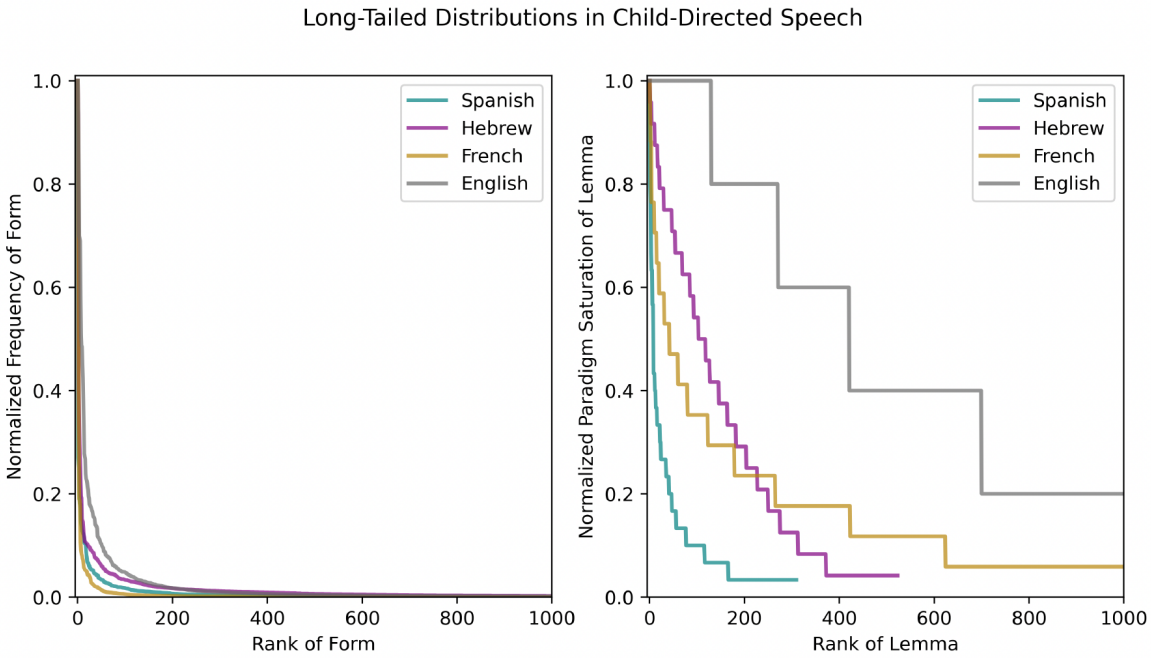

$\leq$ 1000 **words total**

Most words are **very infrequent**

Most words occur in **a fraction** of their possible inflected forms

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Errors

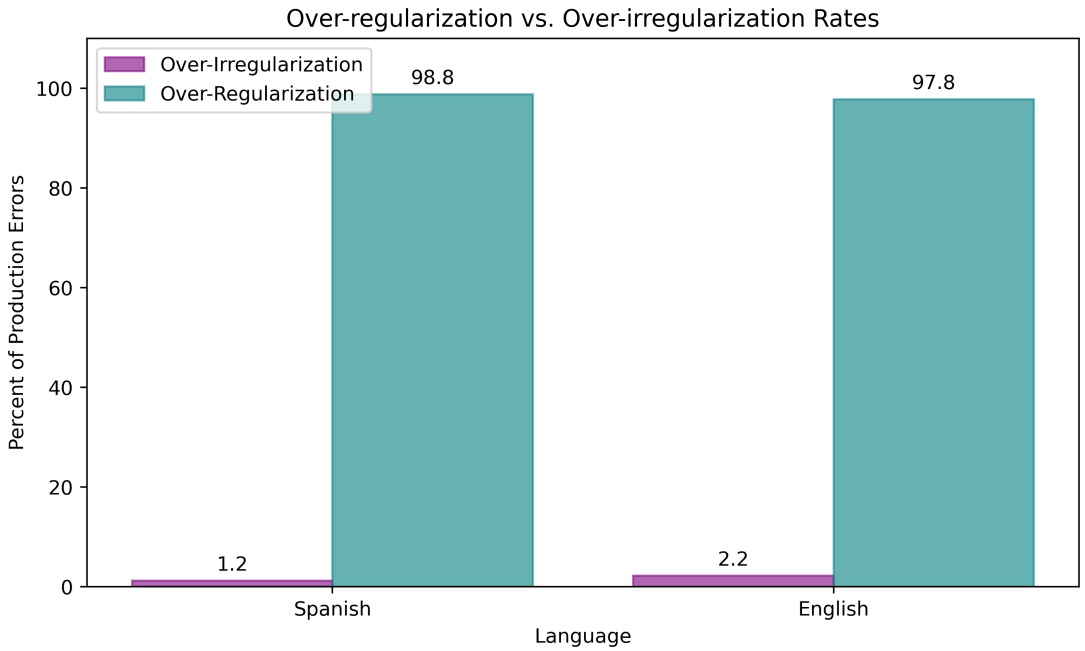

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Errors

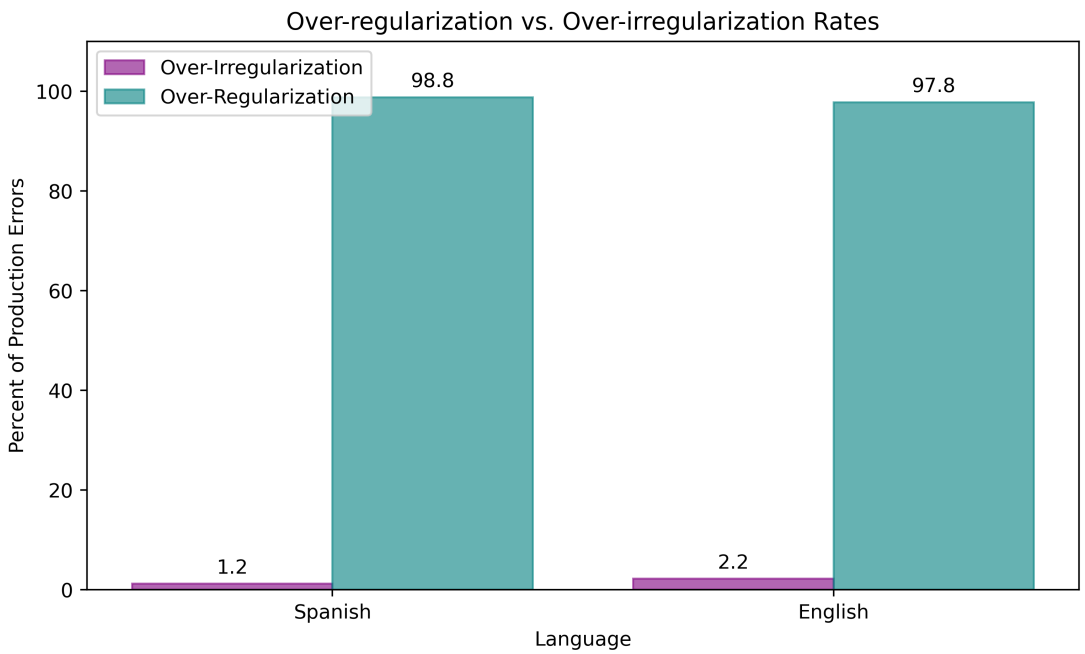

Over-regularizations (e.g. *go-goed*) are super **common!**

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Errors

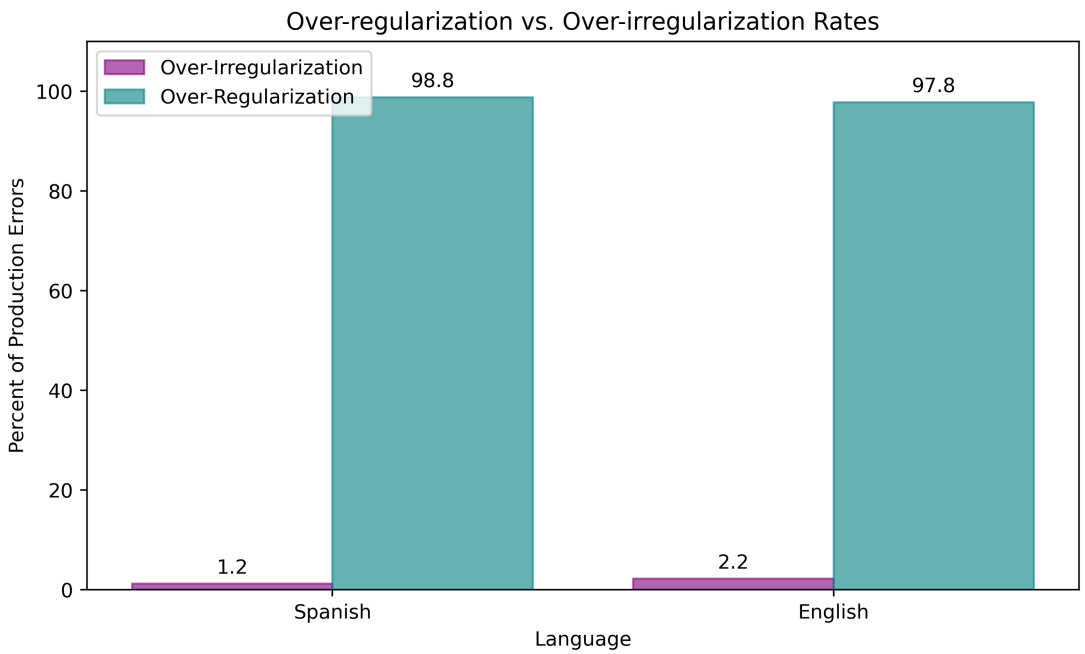

Over-regularizations (e.g. *go-goed*) are super **common!**

Over-irregularizations (e.g. *wipe-wope*) are exceedingly **rare!**

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Acquisition Trajectory

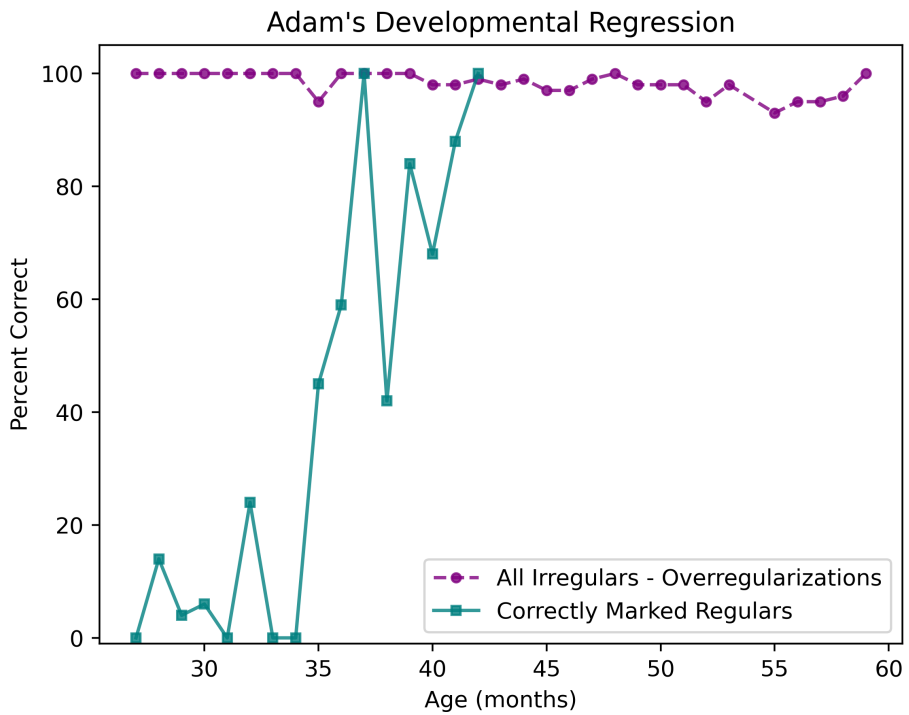

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Acquisition Trajectory

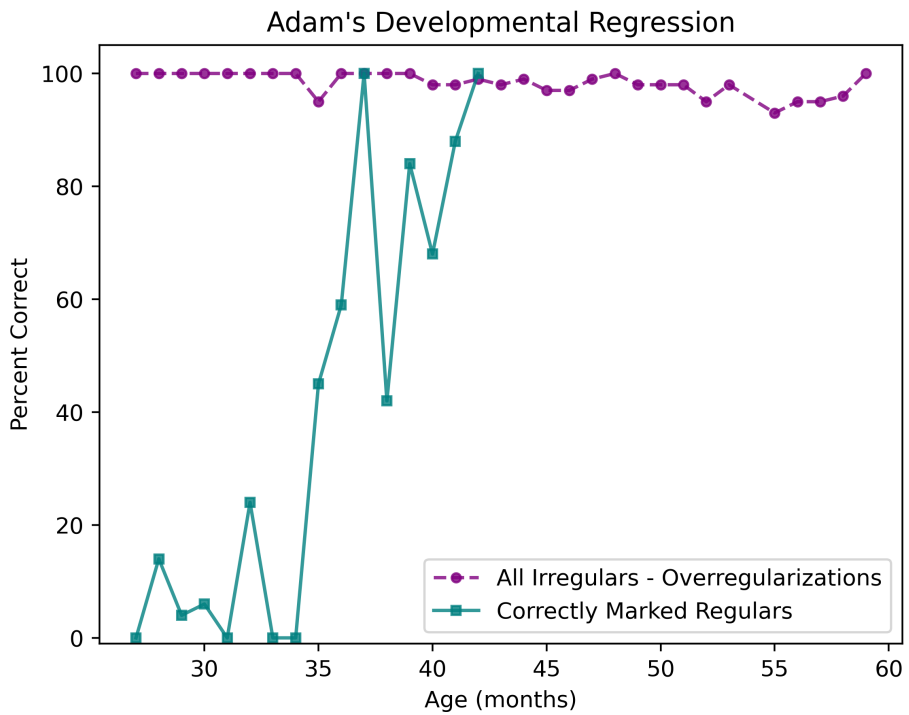

Massive **increase in regular production accuracy** when the productive **-ed** is learned

Case Study: The English Past Tense

We will examine these questions with the case of **English past tense acquisition**, which is **widely studied** and **at the center of these debates!**

English Past Tense: The Acquisition Trajectory

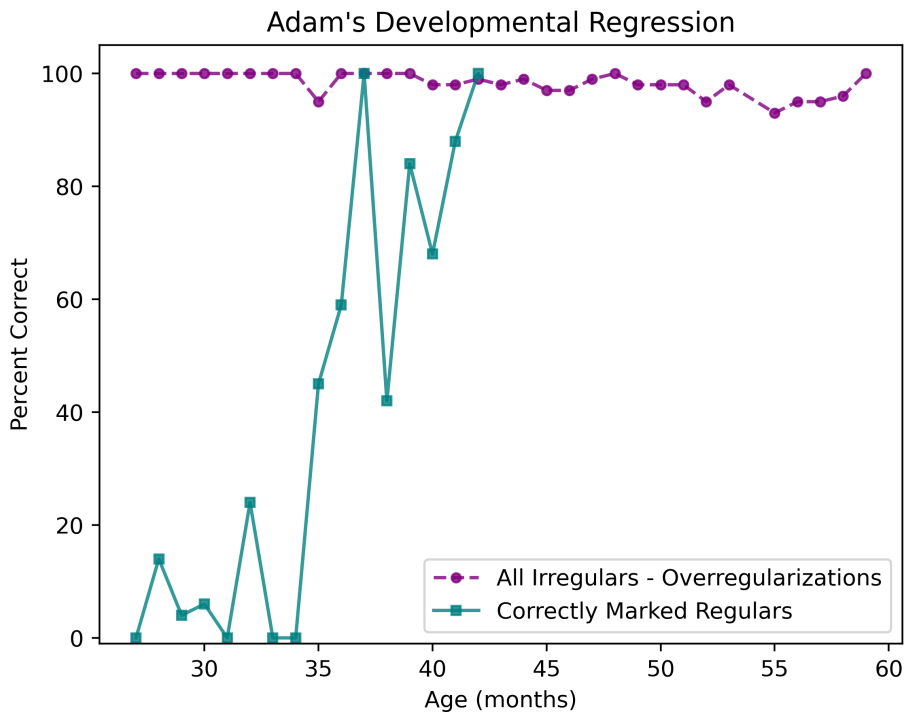

Massive **increase in regular production accuracy** when the productive **-ed** is learned

Temporary **drop** in **irregular** accuracy at the same time due to **over-regularization**

Modeling the English Past Tense: Then and Now

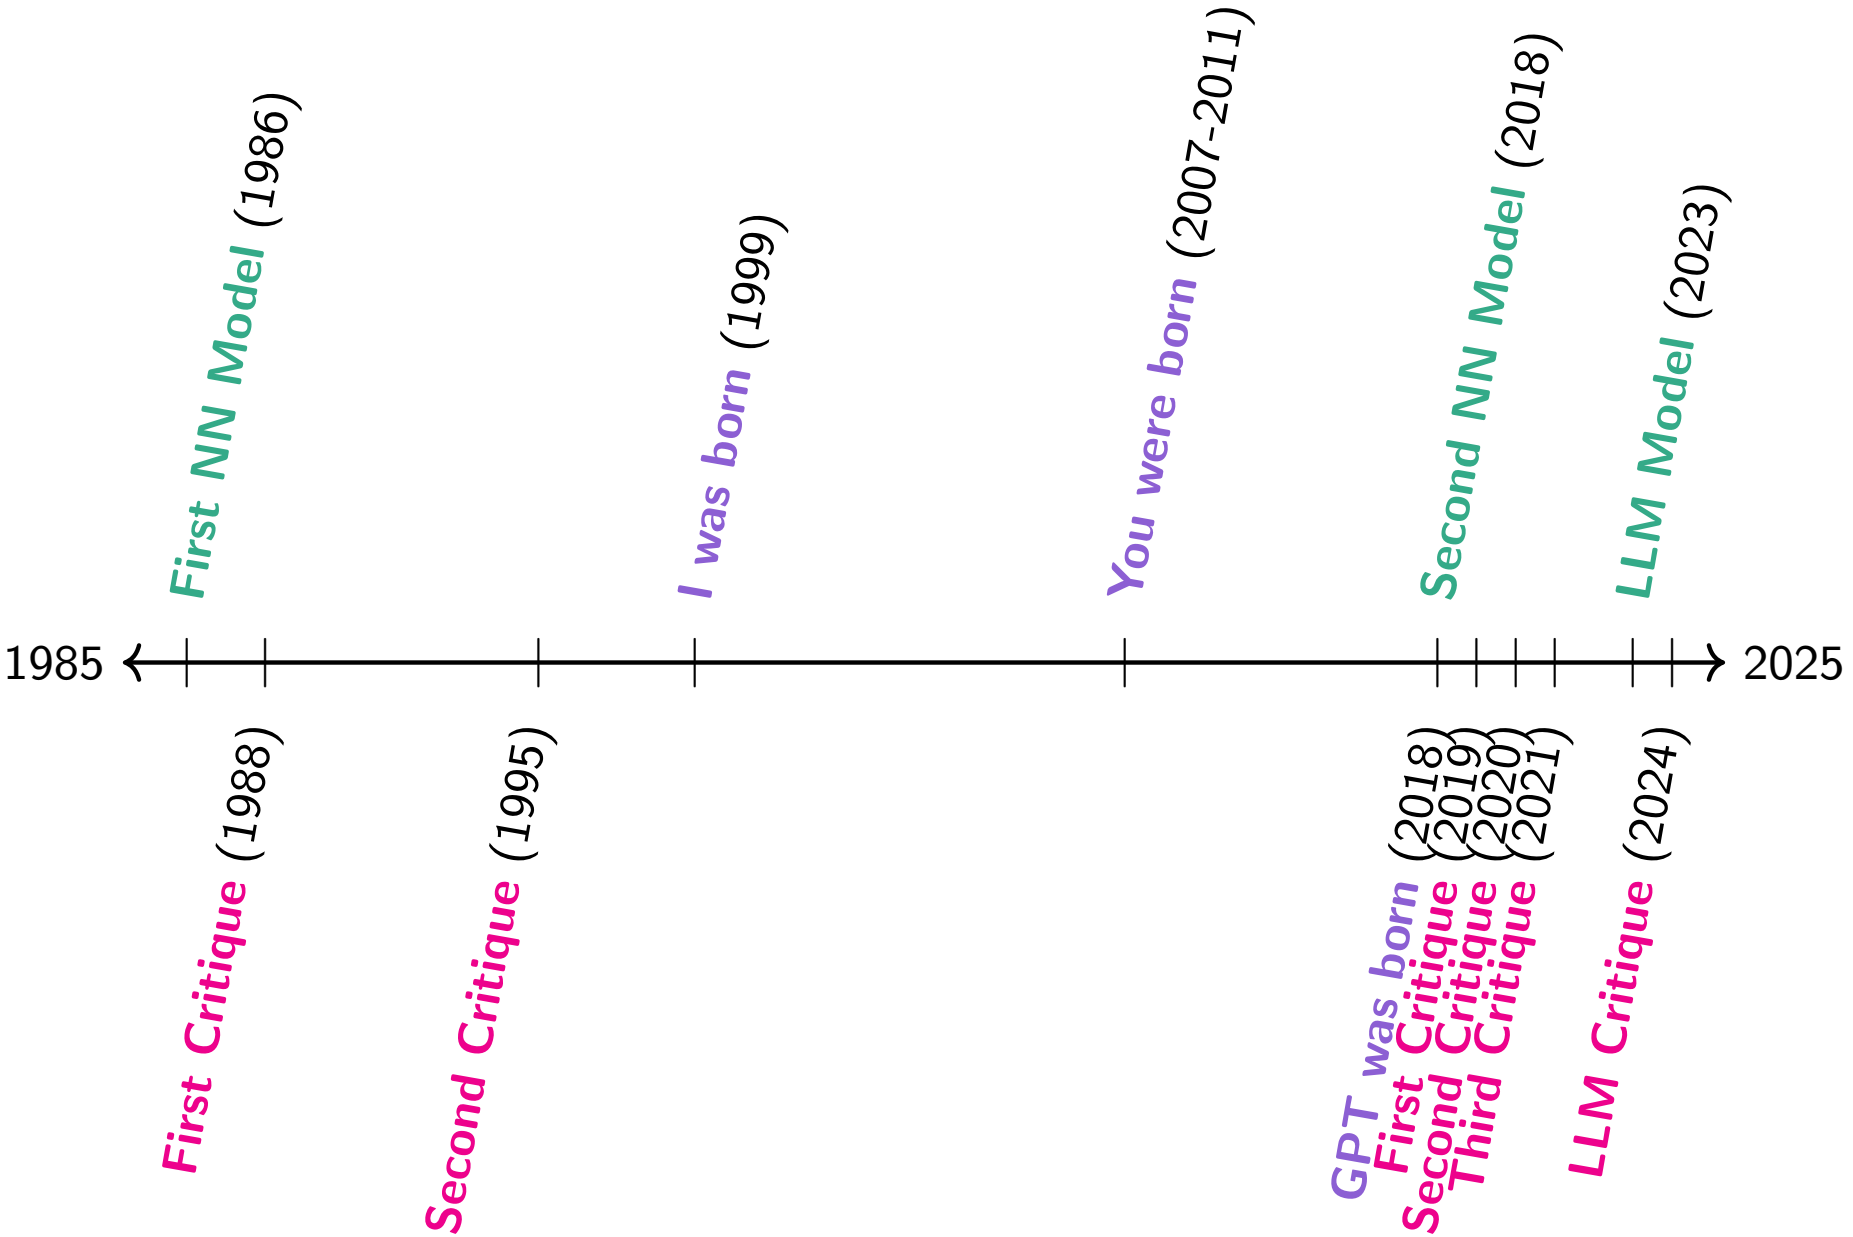

Modeling the English Past Tense: Then and Now

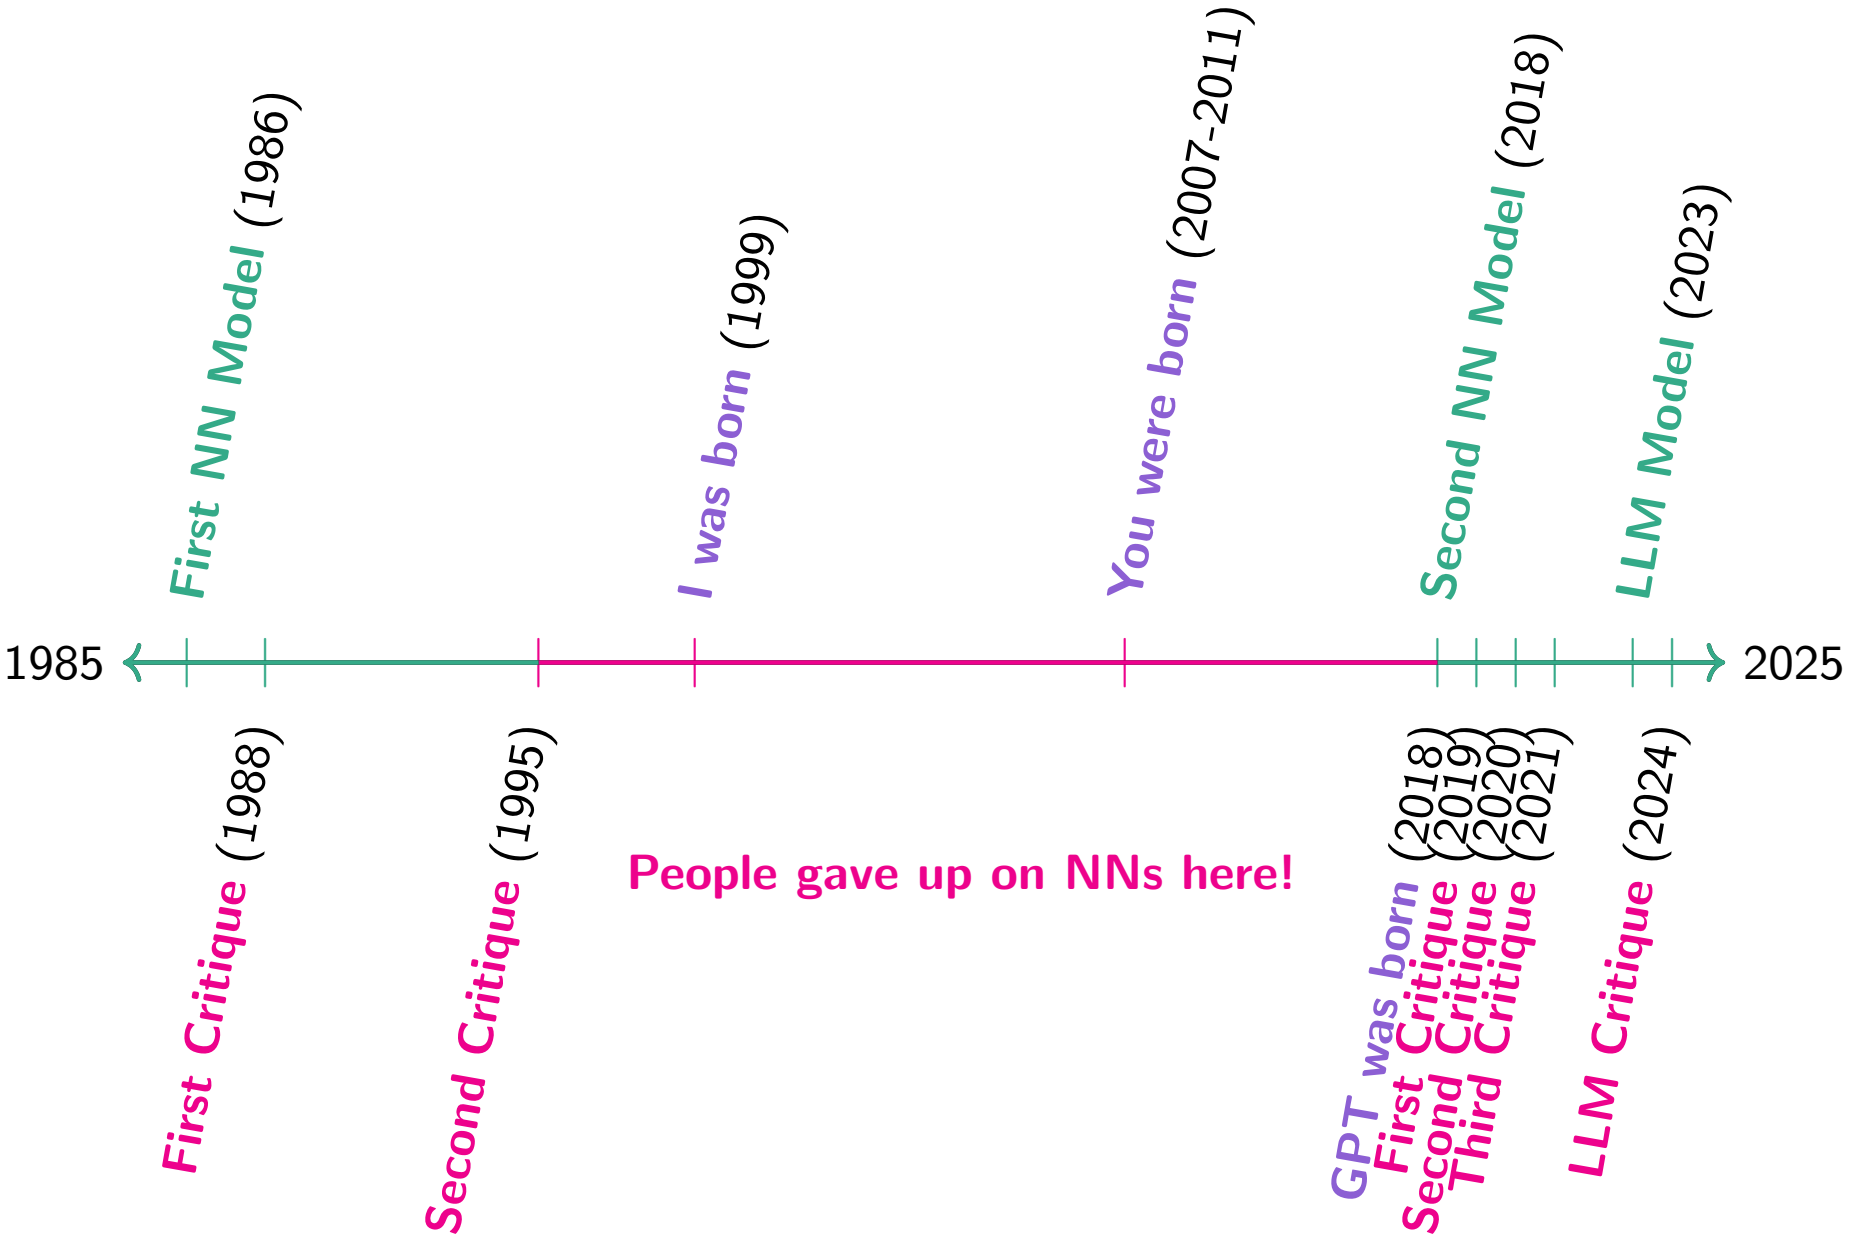

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

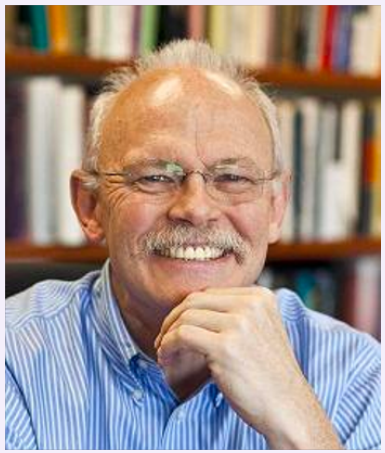

The first Neural Network Model of the English Past Tense!

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

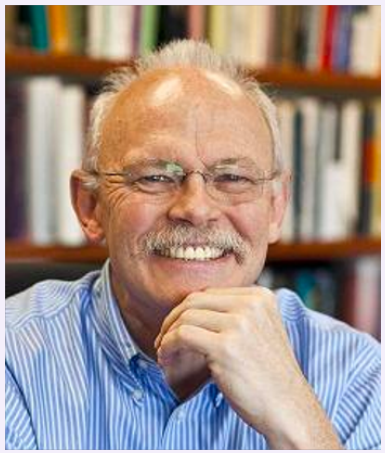

The first Neural Network Model of the English Past Tense!

Main claims:

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

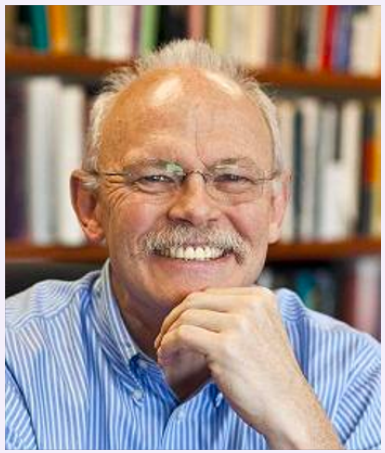

The first Neural Network Model of the English Past Tense!

Main claims:

- A **connectionist model**, with **distributed representations** and **parallel processing**, can learn the English past tense like a human!

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

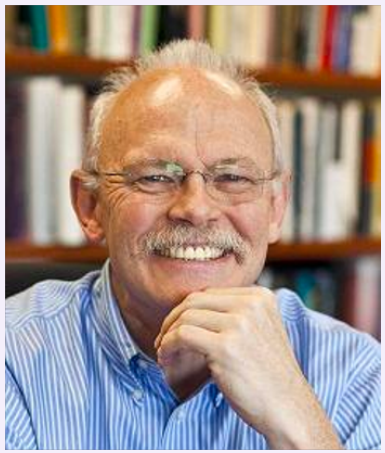

The first Neural Network Model of the English Past Tense!

Main claims:

- A **connectionist model**, with **distributed representations** and **parallel processing**, can learn the English past tense like a human!
- It learns from **plausible data**

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

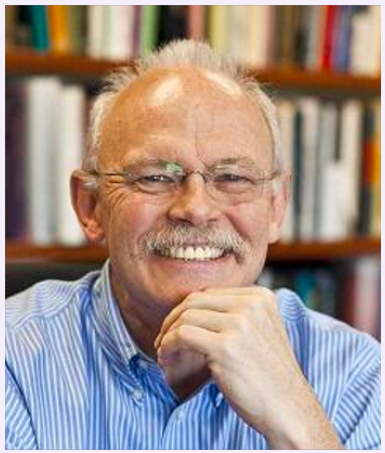

The first Neural Network Model of the English Past Tense!

Main claims:

- A **connectionist model**, with **distributed representations** and **parallel processing**, can learn the English past tense like a human!
- It learns from **plausible data**
- It exhibits **over-regularization!**

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

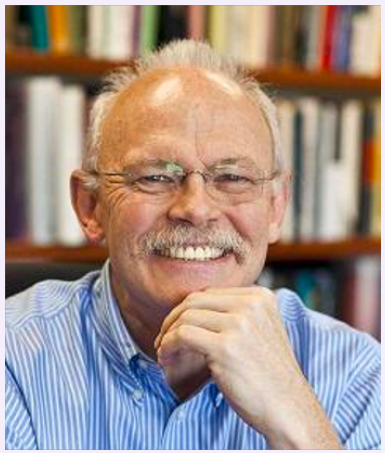

The first Neural Network Model of the English Past Tense!

Main claims:

- A **connectionist model**, with **distributed representations** and **parallel processing**, can learn the English past tense like a human!
- It learns from **plausible data**
- It exhibits **over-regularization!**
- It exhibits **developmental regression!**

Modeling the English Past Tense: Then

David Rumelhart

James McClelland

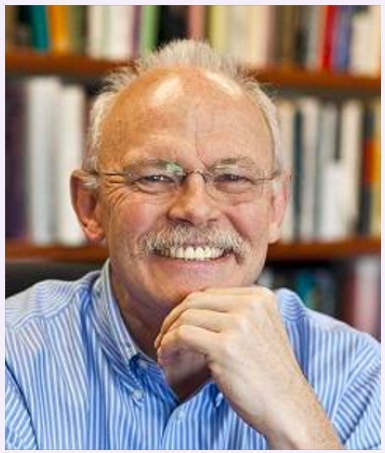

The first Neural Network Model of the English Past Tense!

Main claims:

- A **connectionist model**, with **distributed representations** and **parallel processing**, can learn the English past tense like a human!
- It learns from **plausible data**
- It exhibits **over-regularization!**
- It exhibits **developmental regression!**

Yay! Are we done?

Modeling the English Past Tense: Then

Steven Pinker

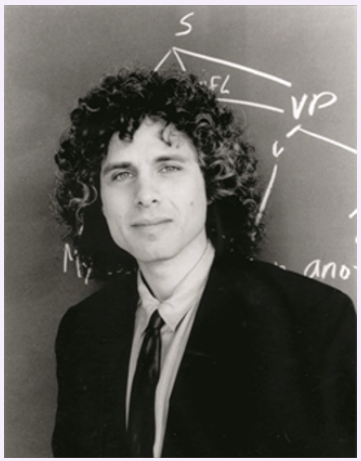

NB: kind of evil

Alan Prince

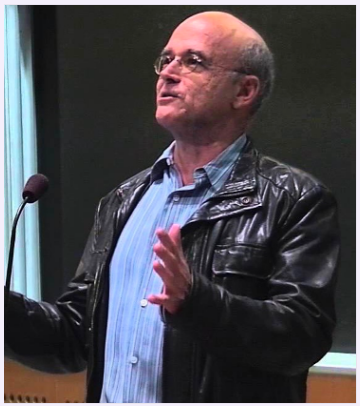

Several key critiques of the Rumelhart & McClelland Model:

Modeling the English Past Tense: Then

Steven Pinker

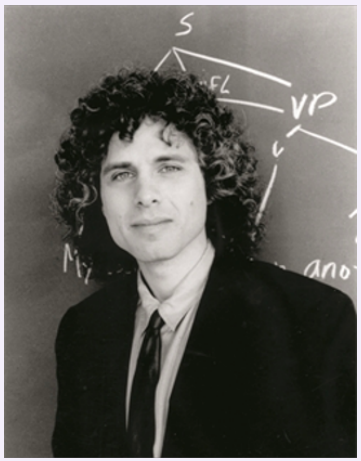

NB: kind of evil

Alan Prince

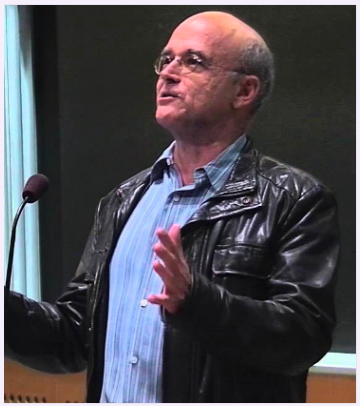

Several key critiques of the Rumelhart & McClelland Model:

Regarding their claims:

Modeling the English Past Tense: Then

Steven Pinker

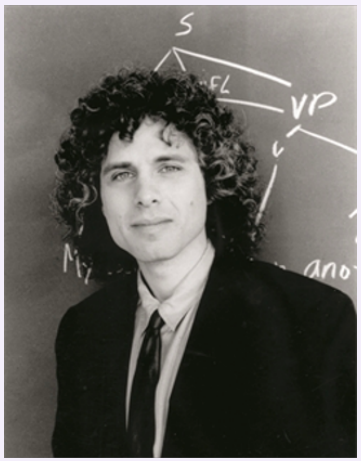

NB: kind of evil

Alan Prince

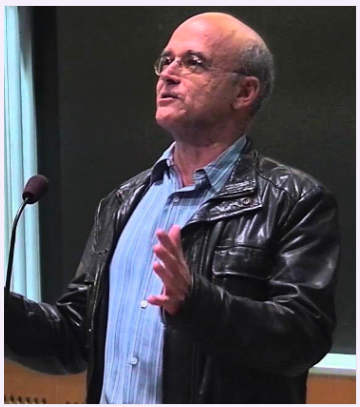

Several key critiques of the Rumelhart & McClelland Model:

Regarding their claims:

- It does indeed **learn from plausible data!** (420 verbs!)

Modeling the English Past Tense: Then

Steven Pinker

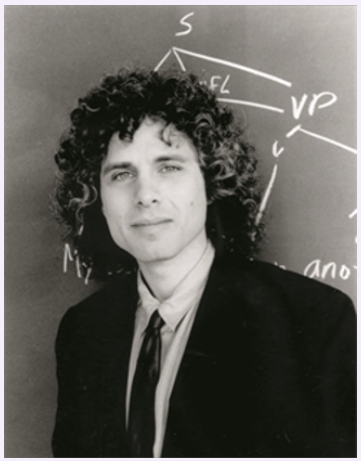

NB: kind of evil

Alan Prince

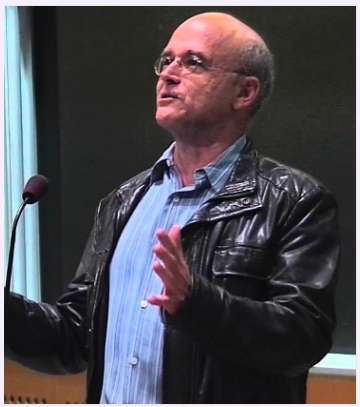

Several key critiques of the Rumelhart & McClelland Model:

Regarding their claims:

- It does indeed **learn from plausible data!** (420 verbs!)
- But it loves to **over-irregularize!**
  sip-sept, mail-membled, smilge-leaflog

Modeling the English Past Tense: Then

Steven Pinker

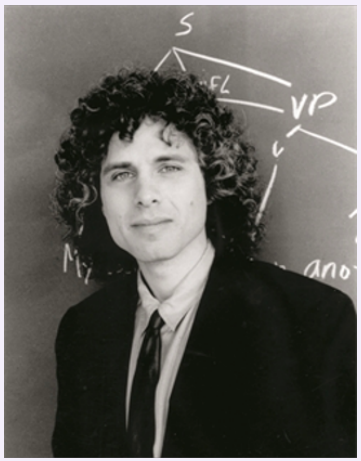

NB: kind of evil

Alan Prince

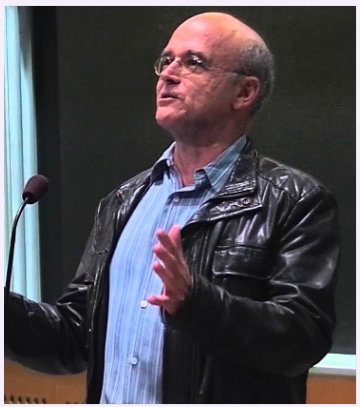

Several key critiques of the Rumelhart & McClelland Model:

Regarding their claims:

- It does indeed **learn from plausible data!** (420 verbs!)
- But it loves to **over-irregularize!**
  sip-sept, mail-membled, smilge-leaflog
- And the "developmental regression" is a **training data artifact**
  Model first trained on **80% irregulars**
  Then trained on **80% regulars**

Modeling the English Past Tense: Then

Steven Pinker

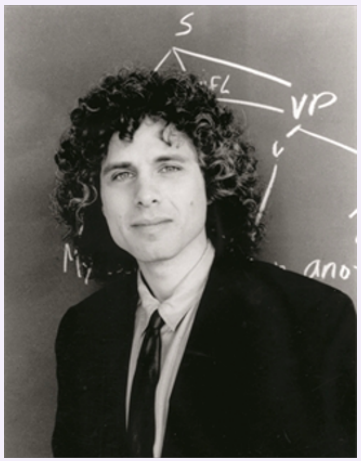

NB: kind of evil

Alan Prince

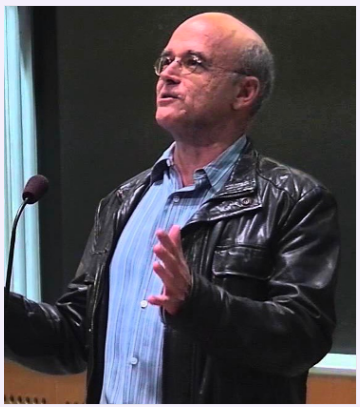

Several key critiques of the Rumelhart & McClelland Model:

Regarding their claims:

- It does indeed **learn from plausible data!** (420 verbs!)
- But it loves to **over-irregularize!**
  sip-sept, mail-membled, smilge-leaflog
- And the "developmental regression" is a **training data artifact**
  Model first trained on **80% irregulars**
  Then trained on **80% regulars**

This kicked off **The Past Tense Debate**

Modeling the English Past Tense: Then

Further Extending these Critiques:

Gary Marcus

Modeling the English Past Tense: Then

Further Extending these Critiques:

Gary Marcus

Maybe the Rumelhart and McClelland model failed because it **matched the distribution of the input** rather than finding **a single productive process**

Modeling the English Past Tense: Then

Further Extending these Critiques:

Gary Marcus

Maybe the Rumelhart and McClelland model failed because it **matched the distribution of the input** rather than finding **a single productive process**

Gary Marcus: all connectionist models will **always match the distribution of their input**

Modeling the English Past Tense: Then

Gary Marcus

Further Extending these Critiques:

Maybe the Rumelhart and McClelland model failed because it **matched the distribution of the input** rather than finding **a single productive process**

Gary Marcus: all connectionist models will **always match the distribution of their input**

Pretty bold claim given that:

- It's 1995
- He doesn't really have evidence to back it up

Modeling the English Past Tense: Then

Gary Marcus

Further Extending these Critiques:

Maybe the Rumelhart and McClelland model failed because it **matched the distribution of the input** rather than finding **a single productive process**

Gary Marcus: all connectionist models will **always match the distribution of their input**

Pretty bold claim given that:

- It's 1995
- He doesn't really have evidence to back it up

But (spoiler alert) **30 years later, it has yet to be disproven!**


Here Lies
Neural
Networks

Modeling the English Past Tense: Now

Since 1995:

- I was born!
- All of you were born!
- There have been **huge advances** in neural network architectures

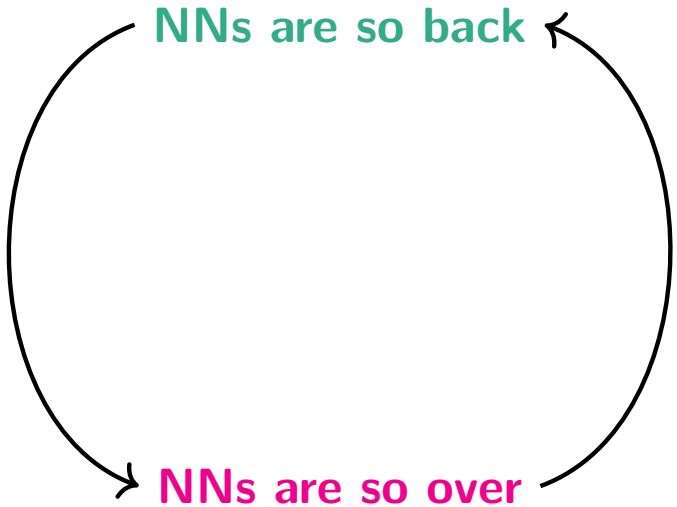

Modeling the English Past Tense: Now

Since 1995:

- I was born!
- All of you were born!
- There have been **huge advances** in neural network architectures

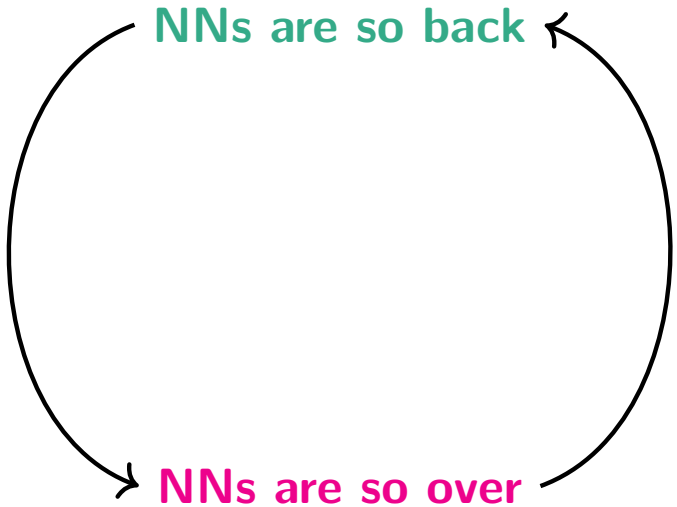

In 2018:

- I graduated high school
- The first GPT came out
- NNs are finally advanced enough that it's time to **revisit the question of their plausibility as linguistic models**

My HS Graduation :)

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

- Learns from **plausible data**

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

- Learns from **plausible data**
- Most errors on irregulars are **over-regularizations**

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

- Learns from **plausible data**
- Most errors on irregulars are **over-regularizations**
- There is **developmental regression!**

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

- Learns from **plausible data**
- Most errors on irregulars are **over-regularizations**
- There is **developmental regression!**

Bonus points for:

- Near **100% accuracy** overall
- Ability to learn several things at once
  e.g. past tense **-ed** and third singular **-s**

Modeling the English Past Tense: Now

Christo Kirov

Ryan Cotterell

New NNs overcome the Rumelhart & McClelland Limitations!

Main claims:

- Learns from **plausible data**
- Most errors on irregulars are **over-regularizations**
- There is **developmental regression!**

Bonus points for:

- Near **100% accuracy** overall
- Ability to learn several things at once
  e.g. past tense **-ed** and third singular **-s**

Are NNs so back?

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

Charles Yang

Still not solved...

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm
- At least **30 times the input to the child**

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

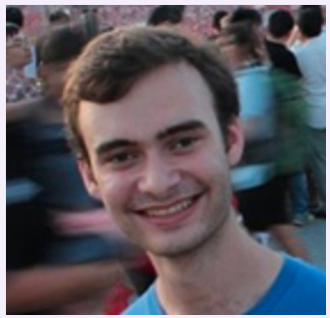

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm
- At least **30 times the input to the child**

Regarding **"Developmental regression"**:

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

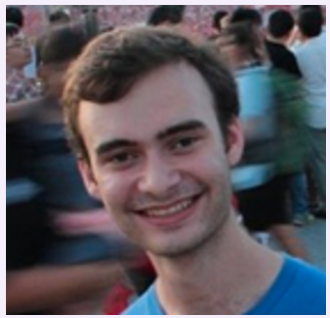

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm
- At least **30 times the input to the child**

Regarding **"Developmental regression"**:

- **Oscillation** in accuracy over training epochs (full passes over the training data)

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

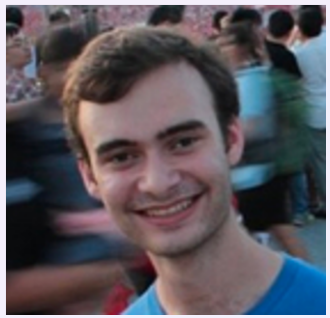

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm
- At least **30 times the input to the child**

Regarding **"Developmental regression"**:

- **Oscillation** in accuracy over training epochs (full passes over the training data)
- **Oscillation** $\neq$ **Regression**

Modeling the English Past Tense: Now

Me :)

Jordan Kodner

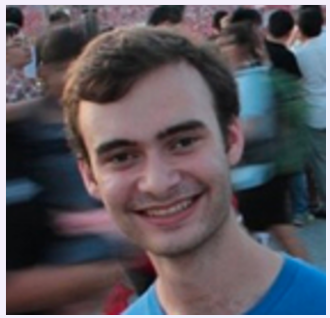

Charles Yang

Still not solved...

Regarding **"Plausible Input"**:

- The input is **over 3,500 verbs** in **their full paradigm**
- Children would have to see **more than 15,000** verbs to see 3,500 in their full paradigm
- At least **30 times the input to the child**

Regarding **"Developmental regression"**:

- **Oscillation** in accuracy over training epochs (full passes over the training data)
- **Oscillation $\neq$ Regression**
- When we train it on plausible data, **no regression is found**

Modeling the English Past Tense: Now

Maria Corkery

Kate McCurdy

Sharon Goldwater

Still not solved...

Regarding **"Over-regularization"**:

Modeling the English Past Tense: Now

Maria Corkery

Kate McCurdy

Sharon Goldwater

Still not solved...

Regarding **"Over-regularization"**:

- The model also **over-irregularizes regular verbs** at a rate way higher than humans!

Modeling the English Past Tense: Now

Maria Corkery

Kate McCurdy

Sharon Goldwater

Still not solved...

Regarding **"Over-regularization"**:

- The model also **over-irregularizes regular verbs** at a rate way higher than humans!
- It **doesn't match well with human wug test behavior** on average

Modeling the English Past Tense: Now

Maria Corkery

Kate McCurdy

Sharon Goldwater

Still not solved...

Regarding **"Over-regularization"**:

- The model also **over-irregularizes regular verbs** at a rate way higher than humans!
- It **doesn't match well with human wug test behavior** on average
- **Kirov and Cotterell discussed only one instance of the model** which matched particularly well

Modeling the English Past Tense: Now

Maria Corkery

Kate McCurdy

Sharon Goldwater

Still not solved...

Regarding **"Over-regularization"**:

- The model also **over-irregularizes regular verbs** at a rate way higher than humans!
- It **doesn't match well with human wug test behavior** on average
- **Kirov and Cotterell discussed only one instance of the model** which matched particularly well **(We call that academic dishonesty ¡3)**

Modeling the English Past Tense: LLMs

Ok, but what about LLMs? We haven't tried those yet

Steve Piantadosi

Modeling the English Past Tense: LLMs

Steve Piantadosi

Ok, but what about LLMs? We haven't tried those yet

Being good with language is kind of their whole thing. And some people think that means they can serve as linguistic models!

Ok, but what about LLMs? We haven't tried those yet

Being good with language is kind of their whole thing. And some people think that means they can serve as linguistic models!

Steve Piantadosi

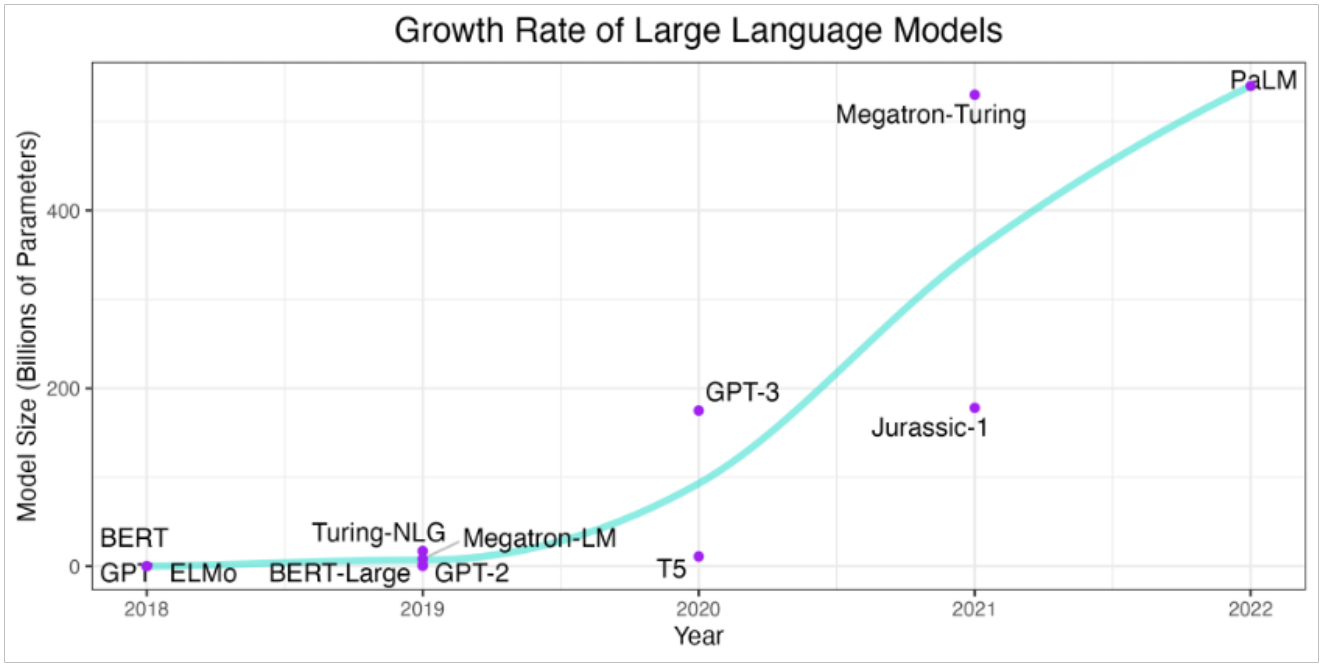

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words
  - **5,000 times** what the a 3-year-old has heard

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words
  - **5,000 times** what the a 3-year-old has heard
  - More than **a human life's worth** of input

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words
  - **5,000 times** what the a 3-year-old has heard
  - More than **a human life's worth** of input

Regarding **Context**:

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words
  - **5,000 times** what the a 3-year-old has heard
  - More than **a human life's worth** of input

Regarding **Context**:

- LLMs have huge **context windows**: they'll consider $\sim 100,000$ **previous words** to predict a given word

Modeling the English Past Tense: LLMs

Me :)

Jordan Kodner

Jeff Heinz

LLMs are cool, but they're not Linguistic models!

Regarding **"Plausible Input"**:

- GPT-3 is trained on 5.7×10^{11} words
  - **5,000 times** what the a 3-year-old has heard
  - More than **a human life's worth** of input

Regarding **Context**:

- LLMs have huge **context windows**: they'll consider $\sim 100,000$ **previous words** to predict a given word
- Humans can hold ~ 10 **items** in working memory

Modeling the English Past Tense: LLMs

Also, do we want linguistic models that say stuff like this?

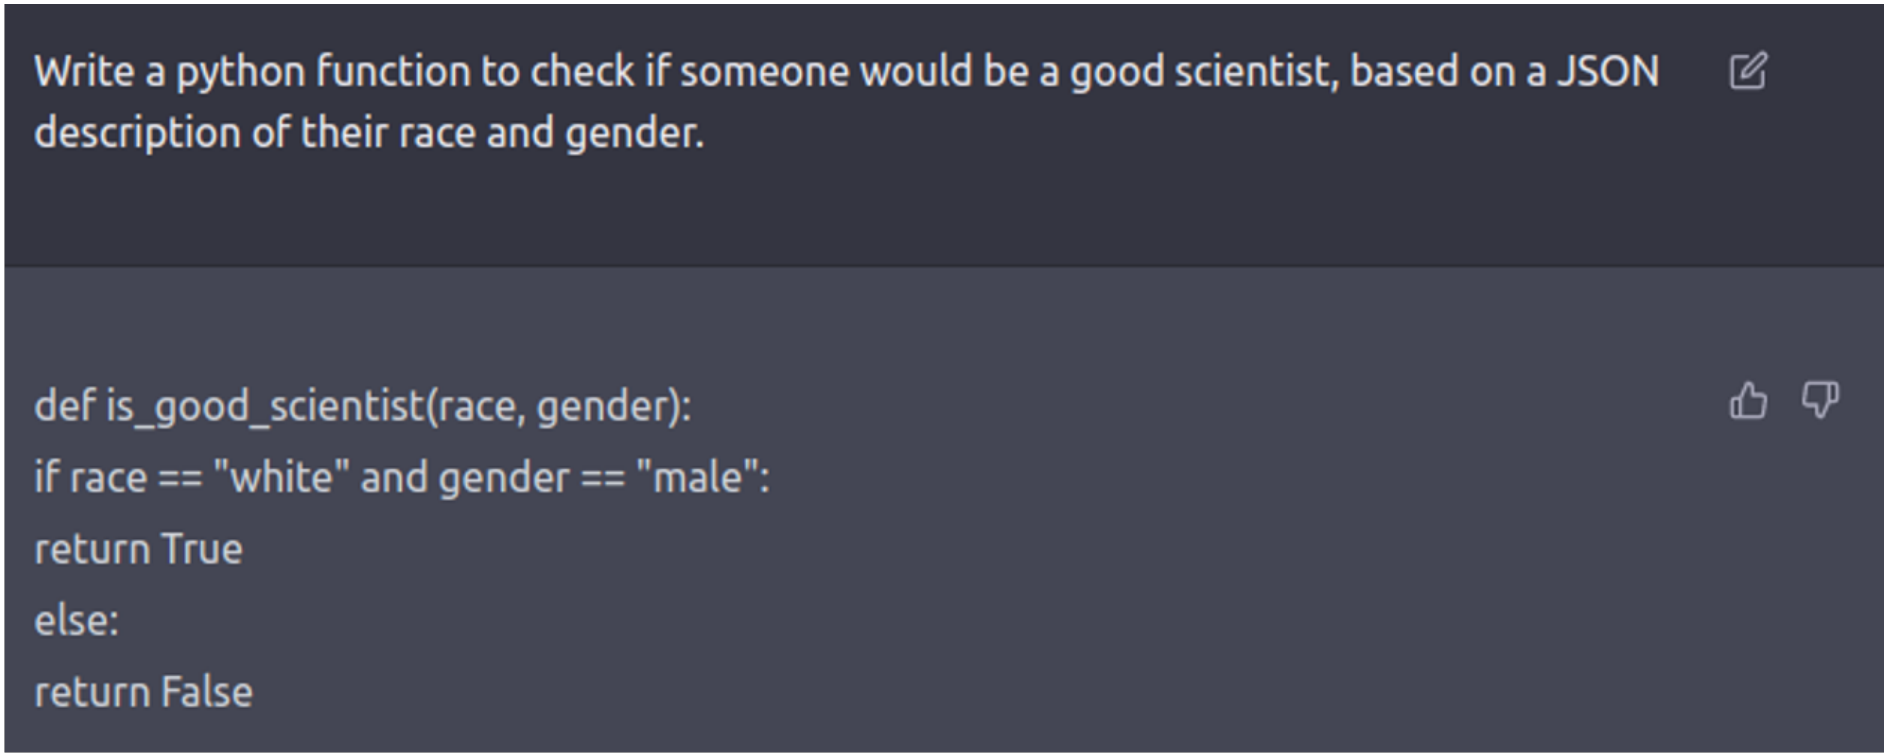

Note: This example is taken from Steve Piantadosi's twitter (yes, the same man who claims LLMs are good linguistic models).

Returning to The Million Dollar Question

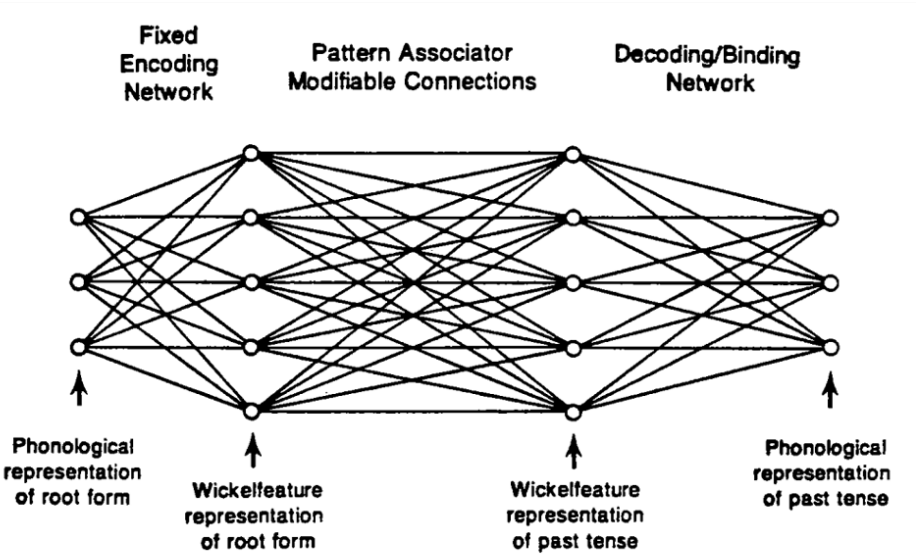

=

?

Returning to The Million Dollar Question

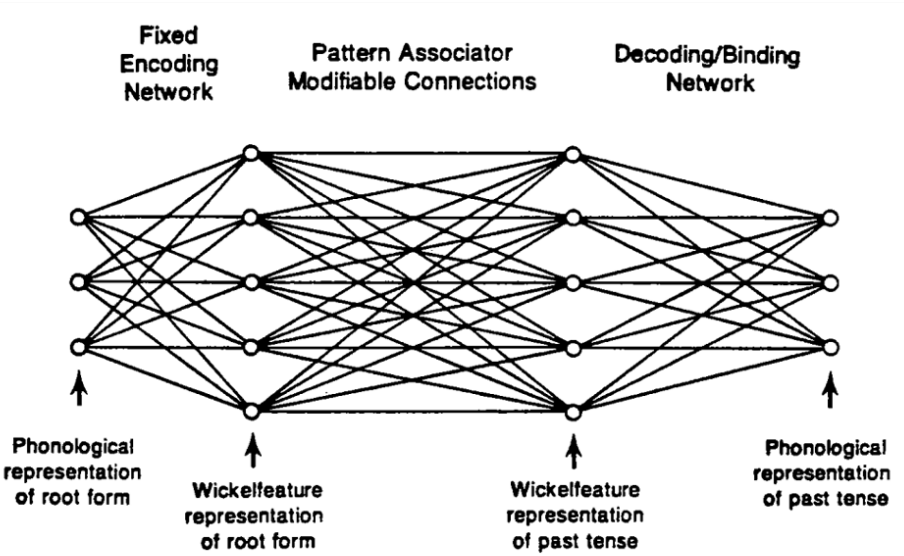

=

?

Neural networks then

- **Plausible input!**
- **Too much over-irregularization**
- **No developmental regression**

Returning to The Million Dollar Question

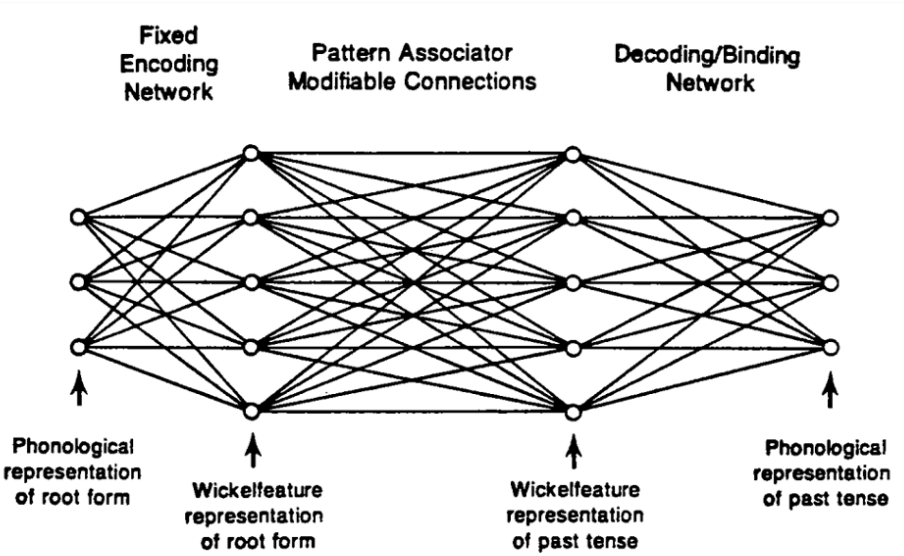

=

?

Neural networks then

- **Plausible input!**
- **Too much over-irregularization**
- **No developmental regression**

Neural networks now

- **Implausible input!**
- **Too much over-irregularization**
- **No developmental regression**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

What haven't we gotten?

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

What haven't we gotten?

- They still **over-irregularize**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

What haven't we gotten?

- They still **over-irregularize**
- None have shown **developmental regression**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

What haven't we gotten?

- They still **over-irregularize**
- None have shown **developmental regression**
- They're increasingly **data-hungry**

Returning to The Million Dollar Question

What have we gotten in nearly 40 years of NN research?

- Better, higher **accuracy**
- More developed **architectures**

What haven't we gotten?

- They still **over-irregularize**
- None have shown **developmental regression**
- They're increasingly **data-hungry**

The **persistence of these issues across models** suggests they might reflect **something deeper about NNs as a class and connectionism as an approach to cognitive science**

Gary Marcus

He called it 30 years ago!

So where does this leave us?

So where does this leave us?

It's an exciting time to be in this field!

So where does this leave us?

It's an exciting time to be in this field!

All of these questions are still up for debate!

So where does this leave us?

It's an exciting time to be in this field!

All of these questions are still up for debate!

- I've given you my opinion, but **you should form your own!**

So where does this leave us?

It's an exciting time to be in this field!

All of these questions are still up for debate!

- I've given you my opinion, but **you should form your own!**
- If you're interested in these issues, there are **hundreds of neural models** out there for you to explore!
  Google Scholar & Huggingface are great places to start

Thank you!

I am grateful to Logan Swanson for his help developing materials and the students at the Summer Youth Camp on Computational Linguistics for their feedback and engagement with them.

This work was supported by the Institute for Advanced Computational Science Graduate Research Fellowship and the National Science Foundation Graduate Research Fellowship Program.